

Simulations of recoiling black holes: adaptive mesh refinement and radiative transfer

Zakaria Meliani¹, Yosuke Mizuno², Hector Olivares², Oliver Porth², Luciano Rezzolla^{2,3}, and Ziri Younsi²

¹ LUTH, CNRS UMR 8102, Observatoire de Paris, Université Paris Diderot, 92190 Meudon, France

² Institute for Theoretical Physics, Goethe-University, D-60438, Frankfurt am Main, Germany

³ Frankfurt Institute for Advanced Studies, D-60438, Frankfurt am Main, Germany

November 11, 2016

ABSTRACT

Context. In many astrophysical phenomena, and especially in those that involve the high-energy regimes that always accompany the astronomical phenomenology of black holes and neutron stars, physical conditions that are achieved are extreme in terms of speeds, temperatures, and gravitational fields. In such relativistic regimes, numerical calculations are the only tool to accurately model the dynamics of the flows and the transport of radiation in the accreting matter.

Aims. We here continue our effort of modelling the behaviour of matter when it orbits or is accreted onto a generic black hole by developing a new numerical code that employs advanced techniques geared towards solving the equations of general-relativistic hydrodynamics.

Methods. More specifically, the new code employs a number of high-resolution shock-capturing Riemann solvers and reconstruction algorithms, exploiting the enhanced accuracy and the reduced computational cost of adaptive mesh-refinement (AMR) techniques. In addition, the code makes use of sophisticated ray-tracing libraries that, coupled with general-relativistic radiation-transfer calculations, allow us to accurately compute the electromagnetic emissions from such accretion flows.

Results. We validate the new code by presenting an extensive series of stationary accretion flows either in spherical or axial symmetry that are performed either in two or three spatial dimensions. In addition, we consider the highly nonlinear scenario of a recoiling black hole produced in the merger of a supermassive black-hole binary interacting with the surrounding circumbinary disc. In this way, we can present for the first time ray-traced images of the shocked fluid and the light curve resulting from consistent general-relativistic radiation-transport calculations from this process.

Conclusions. The work presented here lays the ground for the development of a generic computational infrastructure employing AMR techniques to accurately and self-consistently calculate general-relativistic accretion flows onto compact objects. In addition to the accurate handling of the matter, we provide a self-consistent electromagnetic emission from these scenarios by solving the associated radiative-transfer problem. While magnetic fields are currently excluded from our analysis, the tools presented here can have a number of applications to study accretion flows onto black holes or neutron stars.

Key words. accretion, accretion disks - black hole physics - methods: numerical - radiation: dynamics - relativistic processes

1. Introduction

Many astrophysical phenomena are complex and subject to non-linear dynamics, making numerical simulations an indispensable tool for their study. In the high-energy regimes that always accompany the astronomical phenomenology of compact objects, physical conditions are extreme, with speeds and temperatures so high and gravitational fields so large that both relativistic and general-relativistic effects must be taken into account. In events involving compact objects such as black holes, Einstein's theory of gravitation plays a crucial role, and it is imperative to use it to model accretion flows and radiation therein. In addition to having to model dynamics that are often highly nonlinear, the simulation of compact objects also requires the ability to follow physical phenomena that occur across multiple scales and so must be resolved simultaneously on small and large scales, which requires large amounts of computational resources that cannot be sustained. Adaptive mesh refinement (AMR) provides an effective solution to the problem of performing simulations of phenomena where it is necessary to resolve global as well as local scales.

Over the past few years, great advances in numerical general relativity have given rise to the development of numerical schemes employing the 3 + 1 formulation and Godunov schemes based on approximate Riemann solvers [Rezzolla & Zanotti \(2013\)](#). These advances in numerical general relativity are best described in the reviews by [Font \(2003\)](#) and by [Martí & Müller \(2015\)](#), which provide a thorough description of high-resolution shock-capturing schemes in general-relativistic hydrodynamics (GRHD). Many general-relativistic hydrodynamic and magnetohydrodynamic codes have been developed and evolved over the past three decades ([Hawley et al. 1984](#); [Kudoh 2000](#); [De Villiers & Hawley 2003](#); [Gammie et al. 2003](#); [Baiotti et al. 2005](#); [Duez et al. 2005](#); [Anninos et al. 2005](#); [Antón et al. 2006](#); [Mizuno et al. 2006](#); [Del Zanna et al. 2007](#); [Giacomazzo & Rezzolla 2007](#); [Radice & Rezzolla 2012](#); [Radice et al. 2014](#); [McKinney et al. 2014](#); [Etienne et al. 2015](#); [White & Stone 2015](#); [Zanotti & Dumbser 2015](#)). Some of these implementations provide additional capabilities that incorporate radiation transfer in approximate ways (e.g. [Sądowski et al. 2013](#)) and/or non-ideal magnetohydrodynamics (MHD) regimes (e.g. [Dionysopoulou et al. 2013](#); [Foucart et al. 2016](#)). These codes have been applied to many astrophysical scenarios involving compact objects

and matter. They have been applied to model accretion-ejection, magnetospheres, and compact star structure collapse (e.g. [Dibi et al. 2012](#); [Fragile et al. 2014](#); [McKinney et al. 2014](#)).

In some astrophysical scenarios, adequate modelling can become extremely challenging because of the large disparities in the temporal and spatial scales that may arise in the problem of interest. Under these conditions, approaches employing uniform and non-adaptive grids may become less efficient. These limitations can be overcome by using AMR with adequate refinement or coarsening conditions to sufficiently capture features of interest. The ideal AMR implementation is meant to provide high-resolution simulations at much lower computational cost than uniform-grid methods are capable of. Various AMR strategies exist, such as the patch-based blocks used in ASTROBEAR ([Cunningham et al. 2009](#)), or the full-octree implementations employed in RAMSES ([Teyssier 2002](#)). The strategy implemented in the code used in this paper is the block-octree approach ([van der Holst et al. 2008](#)).

In this paper, we focus on GRHD applications, motivated mainly by our own continued efforts in augmenting the wealth of community codes available for astrophysical research. We first discuss the implementation of general relativistic hydrodynamics with a static background metric. This is then followed by the test of our shock-capturing scheme for GRHD using AMR strategies, which constitutes the core component of modern code development. Code tests with static black hole metrics using two coordinate systems, namely, Boyer-Lindquist and Kerr-Schild (KS), are discussed.

Using two-dimensional (2D) and three-dimensional (3D) general-relativistic numerical simulations that incorporate local AMR, we study the dynamics of a torus in orbit around a recoiling black hole (see, e.g. [Rezzolla 2009](#), for an introductory review on recoiling black holes). Such a kick is likely to result from the merger of supermassive binary black hole systems (SMBBHs). We then calculate the electromagnetic emission from these simulations (images and light curves). In addition to being a perfect testbed for its highly nonlinear and out-of-equilibrium dynamics, the study of the interaction of a recoiling black hole with the surrounding matter has a precise astrophysical application. The analysis of the accretion rate and of the resulting electromagnetic counterparts of recoiling SMBBHs is of great scientific interest, as it will enable the prediction of recoiling signatures when signals from these sources will be detected by the planned space-borne gravitational-wave detector eLISA ([Amaro-Seoane et al. 2012](#)). This is indeed a well-explored area of research, and several studies have investigated the 2D dynamics resulting from a recoiling black hole (e.g. [Corrales et al. 2010](#); [Zanotti et al. 2010](#); [Zanotti 2012](#)) and 3D simulations (e.g. [Lippai et al. 2008](#); [Megevand et al. 2009](#); [Anderson et al. 2010](#); [Ponce et al. 2012](#)). We assess the performance and accuracy of local AMR to perform long-term recoiling black hole simulations within a reasonable amount of computational time.

The structure of the paper is as follows: In Sect. 2 we describe the governing equations, numerical methods for their solution, and numerical test simulations. In Sect. 3 the results of 2D and 3D GRHD simulations of recoiling black holes are presented. In Sect. 4 we describe the general-relativistic radiative transfer formulation and underlying radiative emission model and apply this to the GRHD simulations of recoiling black holes described in Sect. 3. In Sect. 5 we present our conclusions.

Throughout this paper, we use units where the speed of light, $c = 1$, the gravitational constant, $G = 1$, and gas mass is normalised to the central compact object mass. Greek indices run over space and time, that is, $(0, 1, 2, 3)$, and Roman indices run

over space only, that is, $(1, 2, \text{and } 3)$. We assume a signature $(-, +, +, +)$ for the space-time metric. Self-gravity arising from the gas is neglected, and all simulations presented here are made using polar spherical coordinates even though the code also allows for other choices of coordinates.

2. Numerical methods and benchmarks

2.1. GRHD equations and numerical methods

We adopted the 3+1 spacetime decomposition (see, e.g. [Rezzolla & Zanotti 2013](#)), where the metric is given by the line element with the following form,

$$ds^2 = -\alpha^2 dt^2 + \gamma_{ij} (dx^i + \beta^i dt) (dx^j + \beta^j dt), \quad (1)$$

where α is the lapse function, β^i is the shift vector, and γ_{ij} is the three-metric on space-like hypersurface of constant time t . In the $3 + 1$ split of space-time, the metric determinant of space-time $g = \det(g_{\mu\nu})$ relates to the determinant of the purely spatial three-metric as $\gamma = \det(\gamma_{ij}) = -g/\alpha^2$, and only γ is required in what follows.

A perfect non-magnetised fluid is described by four physical variables: the rest-mass density ρ , the thermal pressure p , the specific enthalpy h , and the coordinate-frame four-velocity of the fluid u^μ . With these variables, we can characterise the fluid through the energy-momentum tensor ([Rezzolla & Zanotti 2013](#))

$$T_{\mu\nu} = \rho h u_\mu u_\nu + p g_{\mu\nu}, \quad (2)$$

and an equation of state (EOS), relating the pressure to some of the other thermodynamical properties of the fluid. We used a simple ideal-fluid EOS, $p = (\hat{\gamma} - 1)\rho\epsilon$, where ϵ is the specific internal energy, $\hat{\gamma}$ is the adiabatic index, and the specific enthalpy is given by $h(\rho, p) = 1 + \hat{\gamma}/(\hat{\gamma} - 1)p/\rho$ ([Rezzolla & Zanotti 2013](#)). We here used either $\hat{\gamma} = 5/3$ or $\hat{\gamma} = 4/3$ when modelling an ultrarelativistic fluid.

The fluid evolution is described by the conservation of mass and energy-momentum,

$$\nabla_\mu (\rho u^\mu) = 0, \quad (3)$$

$$\nabla_\mu T^{\mu\nu} = 0, \quad (4)$$

which can be written in a form favourable to conservative numerical integration as

$$\frac{\partial U}{\partial t} + \frac{\partial F^i}{\partial x^i} = S. \quad (5)$$

Here, the vector of conserved variables

$$U \equiv \sqrt{\gamma} \begin{bmatrix} D \\ S_j \\ \tau \end{bmatrix}, \quad (6)$$

is composed of the mass-density $D \equiv \rho W$, of the covariant spatial momentum density $S_j \equiv W \rho h u_j$, and of the total energy density $\tau \equiv W^2 \rho h - p - D$, where we have subtracted the mass-density D to improve accuracy in the nonrelativistic regime. The symbol $W \equiv \alpha u^0 = 1/\sqrt{1 - v^i v_i}$ is the Lorentz factor of the fluid as seen by an Eulerian observer moving with four-velocity $n_\mu = (-\alpha, 0_i)$. The fluxes F^i are then given by

$$F^i \equiv \sqrt{\gamma} \alpha \begin{bmatrix} \rho u^i \\ S_j u^i / W + p \delta_j^i \\ \tau u^i / W + p v^i \end{bmatrix}, \quad (7)$$

and the geometric source terms S are written in terms of Christoffel symbols $\Gamma_{\mu\nu}^\delta$

$$S \equiv \sqrt{\gamma} \alpha \begin{bmatrix} 0 \\ T^{\mu\nu} \left(\frac{\partial g_{\nu j}}{\partial x^\mu} - \Gamma_{\mu\nu}^\delta g_{\delta j} \right) \\ \left(T^{\mu 0} \frac{\partial \alpha}{\partial x^\mu} - \alpha T^{\mu\nu} \Gamma_{\mu\nu}^0 \right) \end{bmatrix}, \quad (8)$$

as discussed by Banyuls et al. (1997), for example. Next to the conserved variables U , computation of fluxes requires a set of “primitive” variables

$$P = \begin{bmatrix} \rho \\ v^i \\ p \end{bmatrix}. \quad (9)$$

A well-known problem of any conservative formulation of the GRHD equations is that while the map $P \rightarrow U$ is straightforward, in general the inverse $U \rightarrow P$ follows as solution to a set of nonlinear equations. We here followed the 1D root-finding algorithm discussed in van der Holst et al. (2008) as it has proven a good compromise between speed and robustness. For a thorough discussion of various algorithms, we refer to the works of Noble et al. (2006), Galeazzi et al. (2013), and Hamlin & Newman (2013) for details.

To evaluate the fluxes F^i in Eq. (7), primitive variables are interpolated to cell interfaces by limited reconstruction using one of the various flavours of piecewise linear slope limiters (e.g. Toro 1999), the piecewise polynomial method (PPM) by Colella & Woodward (1984), or the compact stencil third-order reconstruction by Čada & Torrilhon (2009). This yields the left- and right-biased states U^L and U^R . In the following, superscripts L and R always refers to quantities derived from these reconstructed values. The fluxes are obtained from an approximate Riemann solver.

The two currently available choices are 1.) the Rusanov (LF) scheme, which is based on the knowledge of the maximum absolute value of the characteristic waves at the interface in the direction x : $c^x = \max \{ \lambda_+^{x,L}, \text{abs}(\lambda_-^{x,L}), \lambda_+^{x,R}, \text{abs}(\lambda_-^{x,R}) \}$, and 2.) HLL (Harten et al. 1983), which is based on the knowledge of the two fastest characteristic waves propagating in both directions, one to the left with $c_-^x = \min(\lambda_-^{x,L}, \lambda_-^{x,R})$, and the other wave to the right with $c_+^x = \min(\lambda_+^{x,L}, \lambda_+^{x,R})$. The HLL upwind fluid flux function for the variable U is calculated as

$$\hat{F}^x(U) = \begin{cases} F^x(U^L) & ; \quad c_-^x > 0 \\ F^x(U^R) & ; \quad c_+^x < 0 \\ \bar{F}^x(U^L, U^R) & ; \quad \text{otherwise} \end{cases} \quad (10)$$

where

$$\bar{F}^x(U^L, U^R) \equiv \frac{c_+^x F^x(U^L) - c_-^x F^x(U^R) + c_+^x c_-^x (U^R - U^L)}{c_+^x - c_-^x}, \quad (11)$$

and the LF flux is simply

$$\hat{F}^x(U) = \frac{1}{2} (F^x(U^L) + F^x(U^R)) - \frac{1}{2} c^x (U^R - U^L). \quad (12)$$

According to the chosen stencil, the number of the boundary cells (ghost cells) changes: two cells for linear reconstruction

and three for parabolic reconstruction. On all interior boundaries, these ghost cells are filled by copy/prolongation/restriction operations, depending on the refinement level of bounding grid blocks (see Keppens et al. 2012, for details).

The characteristic wave speed is also used to determine the explicit time step obeying the usual Courant-Friedrich-Levy (CFL) conditions, where the characteristic wave speed λ_\pm^i can be written as (Banyuls et al. 1997)

$$\lambda_\pm^i = \alpha \lambda_\pm'^i - \beta^i, \quad (13)$$

where

$$\lambda_\pm'^i = \frac{(1 - c_s^2) v^i \pm \sqrt{c_s^2 (1 - v^2) [(1 - v^2 c_s^2) \gamma^{ii} - (1 - c_s^2) (v^i)^2]}}{(1 - v^2 c_s^2)}, \quad (14)$$

where $c_s = \sqrt{\partial \ln h / \partial \ln \rho}$ represents the local sound speed and $v = (\gamma_{ij} v^i v^j)^{1/2}$ is the modulus of the local Eulerian velocity of the fluid.

As in any fluid simulation, we cannot handle vacuum, therefore we filled the space of the vacuum region such as outside the torus with a low-density “atmosphere”. This atmosphere had a fixed value for the rest-mass density ρ_{atm} chosen to be several orders of magnitude lower than the highest rest-mass density in the initial disc configuration. A low fixed value for the pressure p_{atm} was chosen, and the material was set to be static with Eulerian velocity $v^i = 0$. At every time step, the values of the primitive variables in the cell were set to the atmospheric values when the density or gas pressure in a given cell fell below the threshold value of $f\rho_{\text{atm}}$ or $f p_{\text{atm}}$. The factor f was chosen for each problem.

We used a block-tree AMR structure where a refinement ratio by a factor of 2 was always ensured between any two successive levels. The global time step was taken as the shortest time step computed across all mesh refinements (Keppens et al. 2012). Prolongation and restriction can be used on conservative variables or primitive variables, where we typically chose the latter to ensure that no unphysical state was encountered in the resolution change. Adaptivity was decided following the simple second-order error estimator by Löhner (Löhner 1987) on physical variables to trigger the AMR. This method is also used in AMR codes such as FLASH (Calder et al. 2002), RAM (Zhang & MacFadyen 2006), MPI-AMRVAC (Keppens et al. 2012; Porth et al. 2014), PLUTO (Mignone et al. 2012), and ECHO (Zanotti et al. 2015), which captures features of interest with sufficient accuracy, as we demonstrate in more detail in Sect. 3.2.

As described by Keppens et al. (2012), the code is parallelised through the message-passing interface (MPI) paradigm. We used a Morton Z-order space-filling curve to run through all blocks in the (oct-) tree data structure. Parallel load balancing was then achieved by allocating equal sections of the space-filling curve to the available processors. Strong- and weak-scaling tests of the underlying MPI-AMRVAC toolkit were performed recently by Porth et al. (2014). In particular, excellent weak scaling to over 30000 processors was demonstrated.

2.2. Spherical accretion: the Michel solution

As a first test of the code for the general-relativistic regime, we considered the stationary solution corresponding to a spherically symmetric solution onto a Schwarzschild black hole. This is known as the Michel accretion solution (Michel 1972) and




Fig. 1. First three columns show 1D radial profiles of the rest-mass density, of the radial component of the four-velocity, and of the gas pressure on the equatorial plane in the Michel solution test in Boyer-Lindquist (upper) and Kerr-Schild coordinates (lower). The semianalytic solution is shown in red. This solution serves as initial condition, while the numerical results at $t = 100 M$ are reported with blue circles. The fourth column shows the entropy increase normalised to the specific heat at constant volume, $\Delta s/c_v$, from $t = 0 M$ to $t = 100 M$ using each coordinate system. A smaller entropy change is related to a better preservation of the stationary solution.




Fig. 2. Snapshots of the 3D distribution of the rest-mass density in the Michel accretion test as computed in Boyer-Lindquist coordinates. The solutions are shown at $t = 0 M$ (left) and at $t = 100 M$ (right).

represents the extension to general relativity of the corresponding Newtonian solution by Bondi (1952). The spherical Michel accretion solution is described in a number of works, (see, e.g. Hawley et al. 1984; Rezzolla & Zanotti 2013). The free parameters are the position of the critical radius r_c and the adiabatic index $\hat{\gamma}$. In this test, we chose $r_c = 8 M$ and $\hat{\gamma} = 5/3$ using an ideal EOS. The steady Bondi accretion flow was initialised, and the simulations were run until $t = 100 M$. We simulated a domain spanning from $2.1 M \leq r \leq 10 M$, $\pi/4 \leq \theta$, and $\phi \leq 3\pi/4$ with 200 cells in all directions.

Figure 1 shows 1D radial profiles of the rest-mass density, of the radial component of the four-velocity, of the pressure, and of the entropy increase normalised to the specific heat at constant volume, $\Delta s/c_v$, along the equatorial plane at $t = 0 M$ and at $t = 100 M$; the top panels in Fig. 1 report the numerical solution

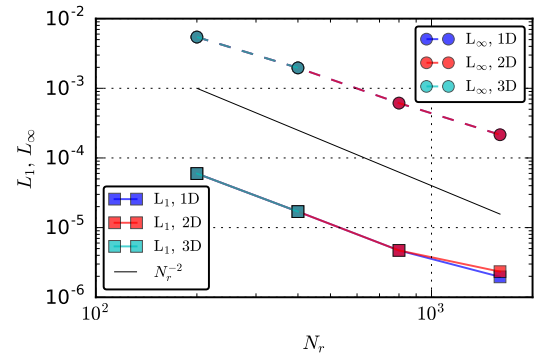


Fig. 3. Convergence order for the Michel accretion test for 1D, 2D, and 3D simulations as computed at $t = 100 M$. Different lines refer either to the dimensionality of the simulations or to the error indicators used, i.e. L_1 or L_∞ norms; we also report as a reference the expected second-order convergence slope.




Fig. 4. 3D isovolume density of a stationary torus at $t = 0 M$ (right panel) and $t = 1500 M$ (left panel) in Kerr space-time using KS coordinates.




Fig. 5. 1D radial profile of density in the equatorial direction of a stationary torus in Kerr space-time using KS coordinates. A resolution of $(N_r, N_\theta) = (200, 100)$ cells was used.

in a space-time expressed in Boyer-Lindquist coordinates, while the bottom panels show it in Kerr-Schild coordinates. The semi-analytic solution is shown in red, while the numerical results are reported with blue circles; we note that although the stationary solution is isentropic, differences in entropy can result from numerical dissipation when an ideal EOS is used.

Clearly, the steady accretion flow is well preserved by the numerical simulations. As is quite common for this test, small differences from the analytic solution are seen near the inner boundary in the case of Boyer-Lindquist coordinates. This occurs because in Boyer-Lindquist coordinates the metric component g_{rr} exhibits severe gradients in the region near the black hole horizon and is therefore difficult to resolve numerically. On the other hand, in Kerr-Schild coordinates, the difference from the analytical solution near the black hole horizon is far smaller, as the Kerr-Schild coordinates are horizon penetrating and lead to a superior behaviour in its vicinity. The 3D structure of the solution is shown in Fig. 2 with a volume-rendering representation of the rest-mass density. The difference between initial and final simulation times is not recognisable, and spherical symmetry is well preserved.

To investigate the numerical accuracy, we checked the L_1 and L_∞ norms in density between initial ($t = 0$) and final simulation time ($t = 100M$) with different resolutions in 1D, 2D, and 3D, as seen in Fig. 3. The convergence is second order for all cases. It is limited by the order of the chosen reconstruction scheme (here we used a Koren slope limiter function).

2.3. Stationary tori

Before turning to the application of a black hole recoiling in a torus, we first verify how well the code is able to preserve a stationary torus solution. In the regime of small kick velocities, it is of particular importance to ensure that the evolution is not governed by numerical artefacts. The hydrodynamic stationary torus solution was first presented in Fishbone & Moncrief (1976) and Kozłowski et al. (1978), and is now a standard test, as used for example by Font & Daigne (2002), Zanotti et al. (2003) and by Antón et al. (2006). In particular, following Font & Daigne (2002), we adopted a non-accreting solution that fills its entire Roche lobe, $\Delta W = 0$ and constant specific angular momentum $\ell = 4.35$ around a Kerr black hole with dimensionless spin parameter $a = 0.5$. We simulated the evolution for ten orbital periods, as measured at the inner torus radius $r_{\text{in}} = 9.34M$. An




Fig. 6. 2D logarithmic rest-mass density of a stationary torus at $t=1500M$ in Kerr space-time using KS coordinates with three different AMR levels.




Fig. 7. L_1 norm of the error in the stationary accretion torus in 2D, with either a uniform grid or with AMR. The error is measured as the difference between the solution at the latest time and the initial data. The red star marks the norm for a test with forced refinement jump in the centre of the torus. All curves indicate second-order convergence, and the AMR cases compare favourably with their corresponding uniform realisations.

orbital period at r_{in} corresponds to $P \approx 150$ time units, and the outer radius of the torus is at $r_{\text{out}} \approx 40M$.

In the following we focus on Kerr-Schild coordinates. The domain covers $r \in [1.85M, 40M]$, which extends into the outer black hole horizon. For the 2D cases, the entire meridional plane is simulated for $\theta \in [0, \pi]$, and in 3D we restrict ourselves to a section with $\Delta\phi, \Delta\theta = \pi/2$ centred on the equator.

As mentioned in Sect. 2.1, to avoid the presence of vacuum regions outside the torus, we applied floor values for the rest-mass density ($\rho_{\text{atm}} = 10^{-5}\rho_{\text{max}}$) and the gas pressure ($p_{\text{atm}} = p_{\text{min}} = 10^{-3}\rho_{\text{atm}}$), where ρ_{max} is the maximum rest-mass density inside the torus at $t = 0$ (this is the rest-mass density at the “centre” of the torus, which we took to be $\rho_c = 1.38 \times 10^{-10} \text{ g cm}^{-3}$). In practice, for all numerical cells that satisfy $\rho \leq f\rho_{\text{atm}}$ or $p \leq fp_{\text{atm}}$, we set $\rho = \rho_{\text{atm}}$, $p = p_{\text{atm}}$, and $v^i = 0$. The factor f is 3.0 for the 2D case and 10.0 for the 3D case.

Volume renderings of density at the initial state and after about ten orbital periods are shown for a resolution of 200^3 cells in Fig. 4. The torus maintains its integrity, and only a slight change in density is observable at the inner edge. An equato-

rial cut of the solution is shown for different times in Fig. 5. As seen already in the 3D rendering, numerical diffusion leads to a smearing out of the inner edge of the torus, but overall, the equilibrium is well maintained for the simulated time. This result gives us confidence about the next set of simulations of recoiling black holes, in which significant departures from the initial state occur on the shorter timescale of one orbital period, which will be captured accurately by the numerical scheme. Similar conclusions have also been obtained when considering a kicked equilibrium torus in the absence of a black hole (the rest-mass density was preserved, but the azimuthal velocity was set to zero and the pressure gradient was removed by suitably adjusting the specific entropy). When performing this test, we obtained that the initial torus shape is well maintained even after being advected multiple times across the simulation domain. We hence conclude that the numerical effect from a velocity kick alone is negligible compared to the physical effects that are due to a recoiling black hole as we study here.

Another important use of the stationary torus solution is that it has allowed us to perform a few controlled experiments with mesh refinement. In the first experiment we allowed for three mesh refinements and let the code automatically refine with Löhner's error estimator on the fluid-frame density. The resulting density map including the grid-structure is illustrated in Fig. 6 after ten orbital periods. Unsurprisingly, the solution in the torus region is essentially identical to the case where a uniform grid at the corresponding resolution was employed.

The Löhner error estimator in essence measures the smoothness of the solution for a given variable, as it depends on a weighted sum of discretised second derivatives. It has the advantage of being more computationally efficient than other error estimators that may require, for instance, solutions computed at different times or different resolutions.

In the case of the recoiling black hole, we relied on automated refinement based on the Löhner error estimator. In a dynamical situation, it necessarily leads to resolution jumps inside the torus. To check how this affects the solution, we performed a second experiment where we enforced the refinement of a single refinement level at one point only, namely, at $(r, \theta) = (15 M, \pi/2)$, which essentially corresponds to the centre of the torus. For all intents and purposes, the result is identical to the case without a resolution jump. The global L_1 error of this setup is illustrated in Fig. 7 for the uniform case, the AMR, and the case with a refinement jump (red star symbol). Overall, we obtain a very good qualitative and quantitative agreement with the uniform-grid case and recover the second-order accuracy of the algorithm.

3. Recoiling black hole

3.1. Motivation for recoiling black hole research

As a full demonstration of the code in a scientific application, we performed 2D and 3D GRHD simulations of recoiling black holes colliding with a circumbinary accretion disc.

Most galaxies are expected to contain a central supermassive black hole that experiences a form of “co-evolution”, which is reflected in a rich phenomenology of black-hole-host galaxy correlations (for a recent review, see Kormendy & Ho 2013). Cosmological models predict that galaxies experience several mergers during their evolution (e.g. Haehnelt 1994; Sesana et al. 2004; Volonteri 2007). Following a galactic merger, the two supermassive black holes will be transported to the barycentre through dynamical friction and form a binary with a sep-

aration of $\lesssim 1$ pc (e.g. Milosavljević & Merritt 2001). Gas funnelled into the galactic centre fuels a (truncated) circumbinary accretion disc with dynamics that are increasingly disconnected from the binary black hole pair (Milosavljević & Phinney 2005). Within the truncated disc, the two black holes lose energy through stellar three-body encounters and by the emission of gravitational waves to finally coalesce and form a single black hole. In the final merger, a significant fraction of the mass, up to 10% (e.g. Reisswig et al. 2009; Barausse et al. 2012), is radiated away, and the produced single black hole can experience a sudden kick with a recoil velocity ranging from several hundred km s^{-1} up to $4,000 \text{ km s}^{-1}$ (e.g. Baker et al. 2007; Gonzalez et al. 2007; Campanelli et al. 2007; Koppitz et al. 2007).

Although no direct evidence of the existence of an SMBBHs system has been found so far, there are several circumstantial possibilities in a number of candidates, such as the radio galaxy 0402+379 (Rodriguez et al. 2006), the ultraluminous infrared galaxy NGC6240 (Komossa et al. 2003), and the BL Lac Object OJ287 (Valtonen et al. 2008). More recently, Graham et al. (2015) reported strong periodic optical variability of the quasar PG 1302-102 with an observed period of 5.2 years. When this optical variability period is matched to the orbital period of the SMBBHs, the system would be separated by less than 0.01 parsecs. This means that the system has evolved well into the final parsec scale.

With the recent first detection of gravitational waves from merging stellar-mass black holes (Abbott et al. 2016), the study of SMBBHs is strongly motivated by the expected detection of their gravitational signal by the space-based gravitational wave detectors, such as the planned eLISA detector (Amaro-Seoane et al. 2012). Considerable attention has recently been attracted by the possibility of detecting the electromagnetic signatures of these events (e.g. Komossa 2012; Schnittman 2013). A number of studies have been carried out to investigate the properties of these electromagnetic signatures either during the stages that precede the merger (Palenzuela et al. 2010; Mösta et al. 2010a; Palenzuela et al. 2010; Moesta et al. 2012; Alic et al. 2012), or in post-merger phase. Several authors have considered the interaction between the binary and the surrounding stars and gas (e.g. Armitage & Natarajan 2002; Milosavljević & Phinney 2005; van Meter et al. 2010; Farris et al. 2010; Farris et al. 2011; Farris et al. 2012; Bode et al. 2012; Giacomazzo et al. 2012; Noble et al. 2012; Gold et al. 2014a). Other scenarios that have not involved matter have also been considered. In these cases, the supermassive black hole binary is considered to be in-spiralling in vacuum, but in the presence of an external magnetic field that is anchored to the circumbinary disc. (e.g. Palenzuela et al. 2009; Mösta et al. 2010b). In the post-merger phase, the electromagnetic counterpart is assumed to be mainly due to the radiation from the circumbinary accretion disc, which will contain an imprint of any strong dynamical change produced on the disc by the merger event. There are two main dynamical effects. One is the abrupt reduction of the rest-mass of the binary that is emitted away in gravitational waves amounting to up to 10 % for equal-mass spinning systems (e.g. Reisswig et al. 2009). The second is the recoil velocity of the merged system, resulting in a kick velocity of the resulting black hole with respect to the host galaxy (e.g. Rezzolla 2009). It is clear that these two dynamical effects can significantly affect the dynamics of circumbinary disc, mainly in their contribution to the formation and propagation of shocks, thereby enhancing the possibility of a strong electromagnetic signal. Several authors have discussed the dynamics and related emission from a circumbinary disc with the recoiling central black hole in the post-merger phase (e.g. Lippai et al.

2008; Megevand et al. 2009; Anderson et al. 2010; Corrales et al. 2010; Rossi et al. 2010; Zanotti et al. 2010; Ponce et al. 2012; Zanotti 2012; Gold et al. 2014b).

3.2. Initial setup in 2D

For the initial setup of the recoiling black hole in 2D, we followed the work by Zanotti et al. (2010). As the initial model of the circumbinary disc, we adopted a stationary disc with a density and pressure profile similar to that of the equatorial plane of the torus described in Sect. 2.3. Similarly to Zanotti et al. (2010), we assumed that the vertical structure of the disc can be neglected and the vertical thickness can be approximated by a quantity $2H$ that is constant in the radial direction, as in the standard thin-disc approximation. The spin parameter of the black hole was $a = 0.5$, and the distribution of specific angular momentum ℓ on the equatorial plane was constant, with $\ell = 8$. The inner and outer edges of the torus were located at $r_{\text{in}} = 40 M$ and $r_{\text{out}} \approx 116 M$, respectively, and the EOS of the fluid was that of an ideal gas with adiabatic index $\gamma = 4/3$. The setup therefore corresponds to the model referred to as S. 50 in the work of Zanotti et al. (2010). Instead of adding a recoil velocity to the black hole as described in the previous section, we evolved the system in a frame where the black hole was fixed and performed a Lorentz boost on the fluid velocity to account for the recoil velocity of the black hole. This has the advantage that the black hole remains at the centre of the coordinate system and the metric functions need not be updated. In all cases considered, the magnitude of the recoil velocity was $v_R = 10^{-3}$ and was directed along the positive x axis.

The simulation domain covers $r \in [1.85, 400]$, which extends into the outer black hole horizon since we used Kerr-Schild coordinates. We performed the simulations at the equatorial plane ($\theta = \pi/2$) with $\phi \in [0, 2\pi]$. To test convergence and the efficiency of AMR with respect to uniform runs, we performed three simulations with uniform resolutions, namely $N_r = 256, 512$, and 1024 , and $N_\phi = \frac{1}{2}N_r$, which we term low, medium, and high resolutions, respectively.

We performed four AMR simulations, two using two refinement levels and two using three levels. The base level has the same resolution as the lowest resolution uniform run, and the cell dimensions are halved when a region moves up one level. In this way, the highest level in a 2 (3) AMR level run has a resolution equivalent to that of the medium- (high-) resolution uniform run. As mentioned above, the refinement of the mesh was automated and the decision of refining or coarsening a given block was taken based on the Löhner estimator. The difference between each of the two pairs of simulations with the same number of AMR levels lies in the tolerance prescription. In each pair, one of the runs has a tolerance $\varepsilon_t = 0.1$ and the other has a tolerance $\varepsilon_t = 0.005$. Clearly, having a lower tolerance implies that the AMR is switched on more frequently, so that in the limit of $\varepsilon_t \rightarrow 0$, the highest refinement level is always present.

The evolution was carried out up to $t = 20,000 M$, corresponding to ~ 15 orbital periods. We here also made use of an atmosphere with a fixed value for the rest-mass density ρ_{atm} chosen five orders of magnitude lower than the highest rest-mass density in the initial disc configuration. A low fixed value for the pressure p_{atm} was chosen, and the material was set to be static with Eulerian velocity $v^i = 0$. In every time step, the values of the primitive variables in the cell were set to the atmospheric values when the density in a given cell fell below the threshold value of $f\rho_{\text{atm}}$ with $f = 1.5$.

3.3. Results in 2D

Figure 8 shows the logarithmic density of the fluid at four different simulation times for a three-level AMR simulation and a high-resolution uniform simulation. AMR and uniform grid cases exhibit very similar features. The asymmetry introduced by the kick direction induces an accumulation of gas on one side of the disc, with a corresponding significant decrease in density on the opposite side of the disc. As time progresses, the variation in density and size of the disc increases. Around $t = 10,000 M$, a part of the disc matter accretes onto the central black hole. By that time, the motion of the recoiling black hole in the plane of the accretion disc has induced spiral shocks that move outwards on a timescale that is comparable with the orbital timescale. These shocks expand from the inner parts of the disc and help to transport angular momentum outwards in the later evolutionary stage. It is worth mentioning that even for an AMR simulation with relatively high tolerance, a good qualitative agreement with the equivalent high-resolution run can be seen, despite the presence of strong shocks and complex dynamics.

The accurate determination of the position of the shock is important for studying the dynamics of the recoiling black hole and for a correct calculation of the emitted radiation. In previous studies (Lippai et al. 2008; O'Neill et al. 2009; Megevand et al. 2009), the propagation of a spiral caustic and a possible shock was inferred only by checking the density and/or pressure gradients. Corrales et al. (2010) introduced a more accurate shock detector presented in the FLASH code. However, these methods are rather empirical criteria and cannot be used to detect weak shocks.

To improve the sensitivity in the determination of the shock position, here we used a relativistic shock detector that exploits an idea proposed in Zanotti et al. (2010) (see also Rezzolla & Zanotti 2002; Rezzolla et al. 2003, for more details). It consists of the possibility of predicting the outcome of the wave pattern in a Riemann problem. In brief, given the left and right states of a Riemann problem, it is possible to compute the threshold relative velocity between them that are required to produce a shock. The actual relative velocity between the two states is compared to this value, and when it exceeds it, the region is marked as shocked. The shock location obtained in this way is shown in Fig. 9. The development of a spiral shock in the accretion disc is clearly seen. In the left panels of Fig. 9, we also plot the AMR blocks at higher refinement levels (levels two and three). It is also quite clear from Fig. 9 that the Löhner scheme (Löhner 1987) used for estimating the error and triggering refinement is very effective and triggers a refinement level even when the shock is rather weak (cf. the trailing edge of the spiral shock). Because of its intrinsic simplicity, it may be preferable to the Rezzolla-Zanotti shock detector when the location of the shock is not of paramount importance.

To analyse the effect of AMR on the dynamics of the disc, we calculated the internal energy (i.e. the volume integral of the internal energy density $\rho\epsilon$), which can later be compared to the light curves of the calculated thermal radiation. In Fig. 10 we show the evolution of this quantity using uniform grids with different resolution and using AMR with different tolerances. The overall behaviour is similar for all cases. Initially, very small oscillations are seen. This is related to the epicyclic picture of the density maximum in the disc. The sharp drop occurs when the disc matter is accreted onto the central black hole. Then a sharp rebound occurs as a result of the development of the shock in the disc. In the later evolutionary stage, the internal energy maintains a higher value with small oscillations. A

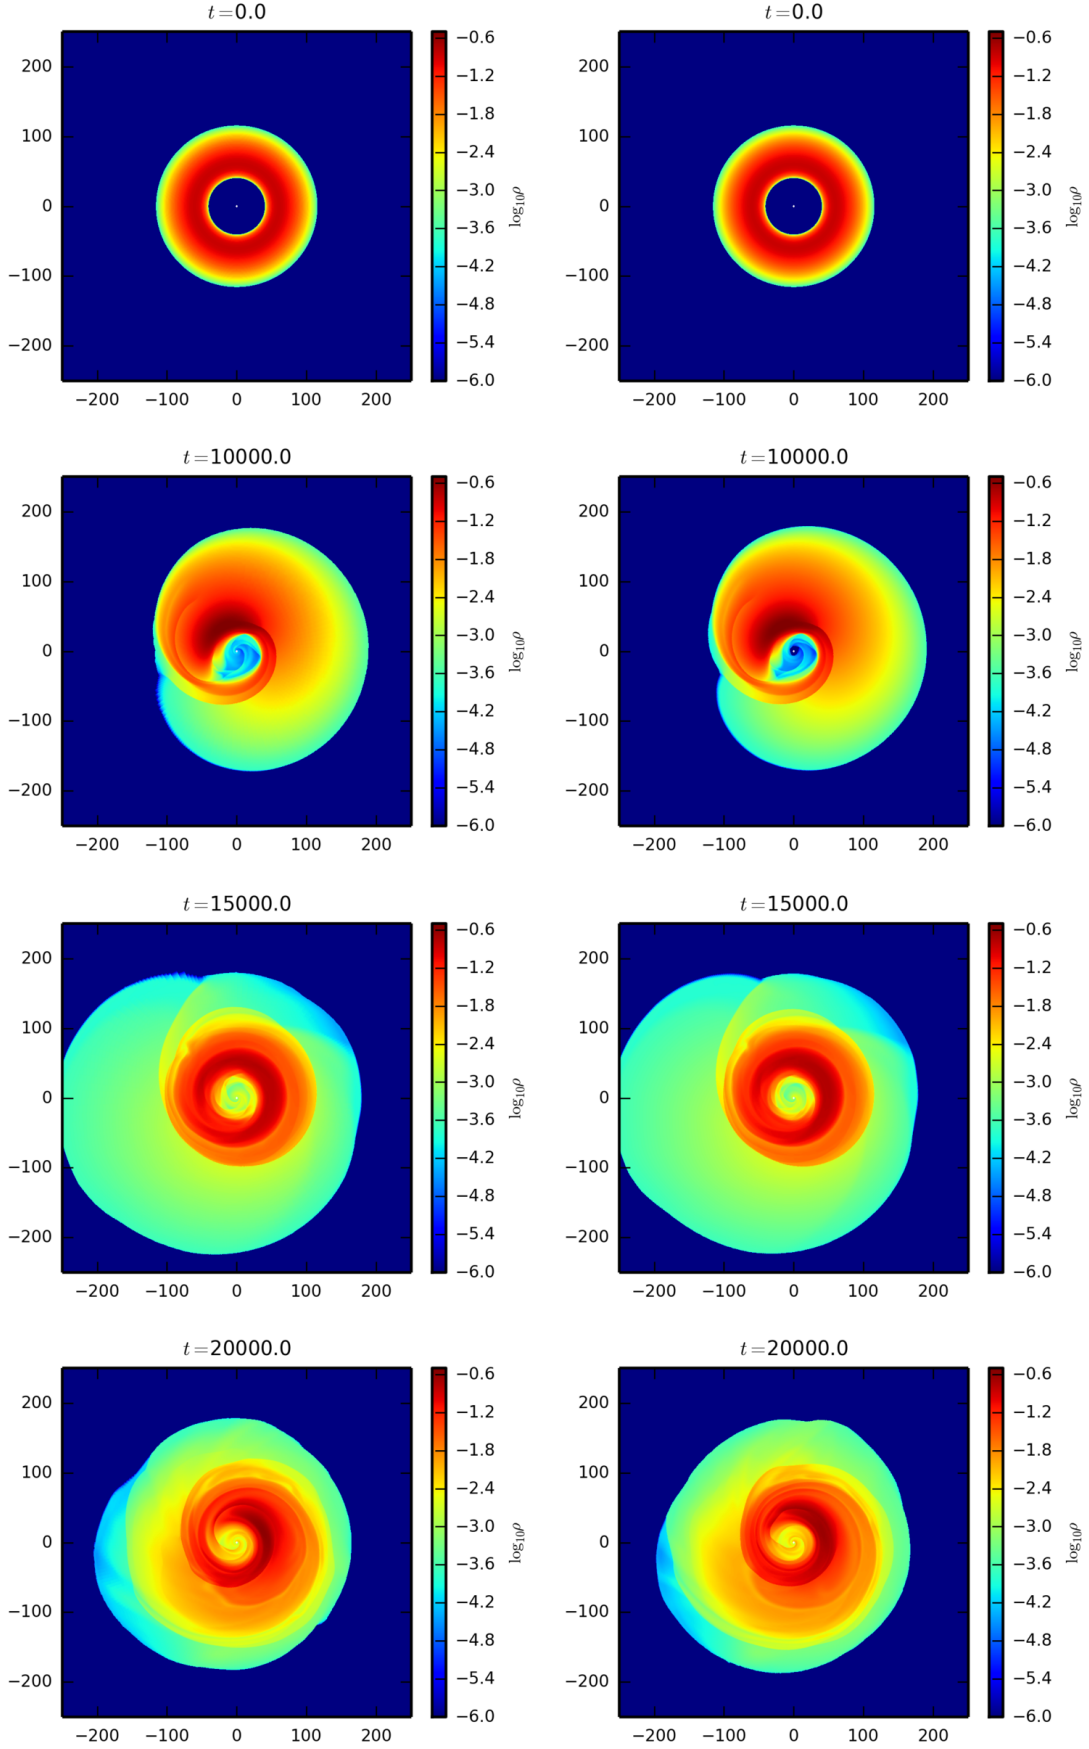


Fig. 8. Evolution of the logarithmic density in the 2D recoiling black hole simulations at times 0, 10,000, 15,000, and 20,000 M , with three AMR levels and tolerance $\varepsilon_t = 0.1$ (left columns) and a high-resolution uniform grid (1024×512 , right columns). At $t = 0 M$, the black hole is moving along the positive x direction.

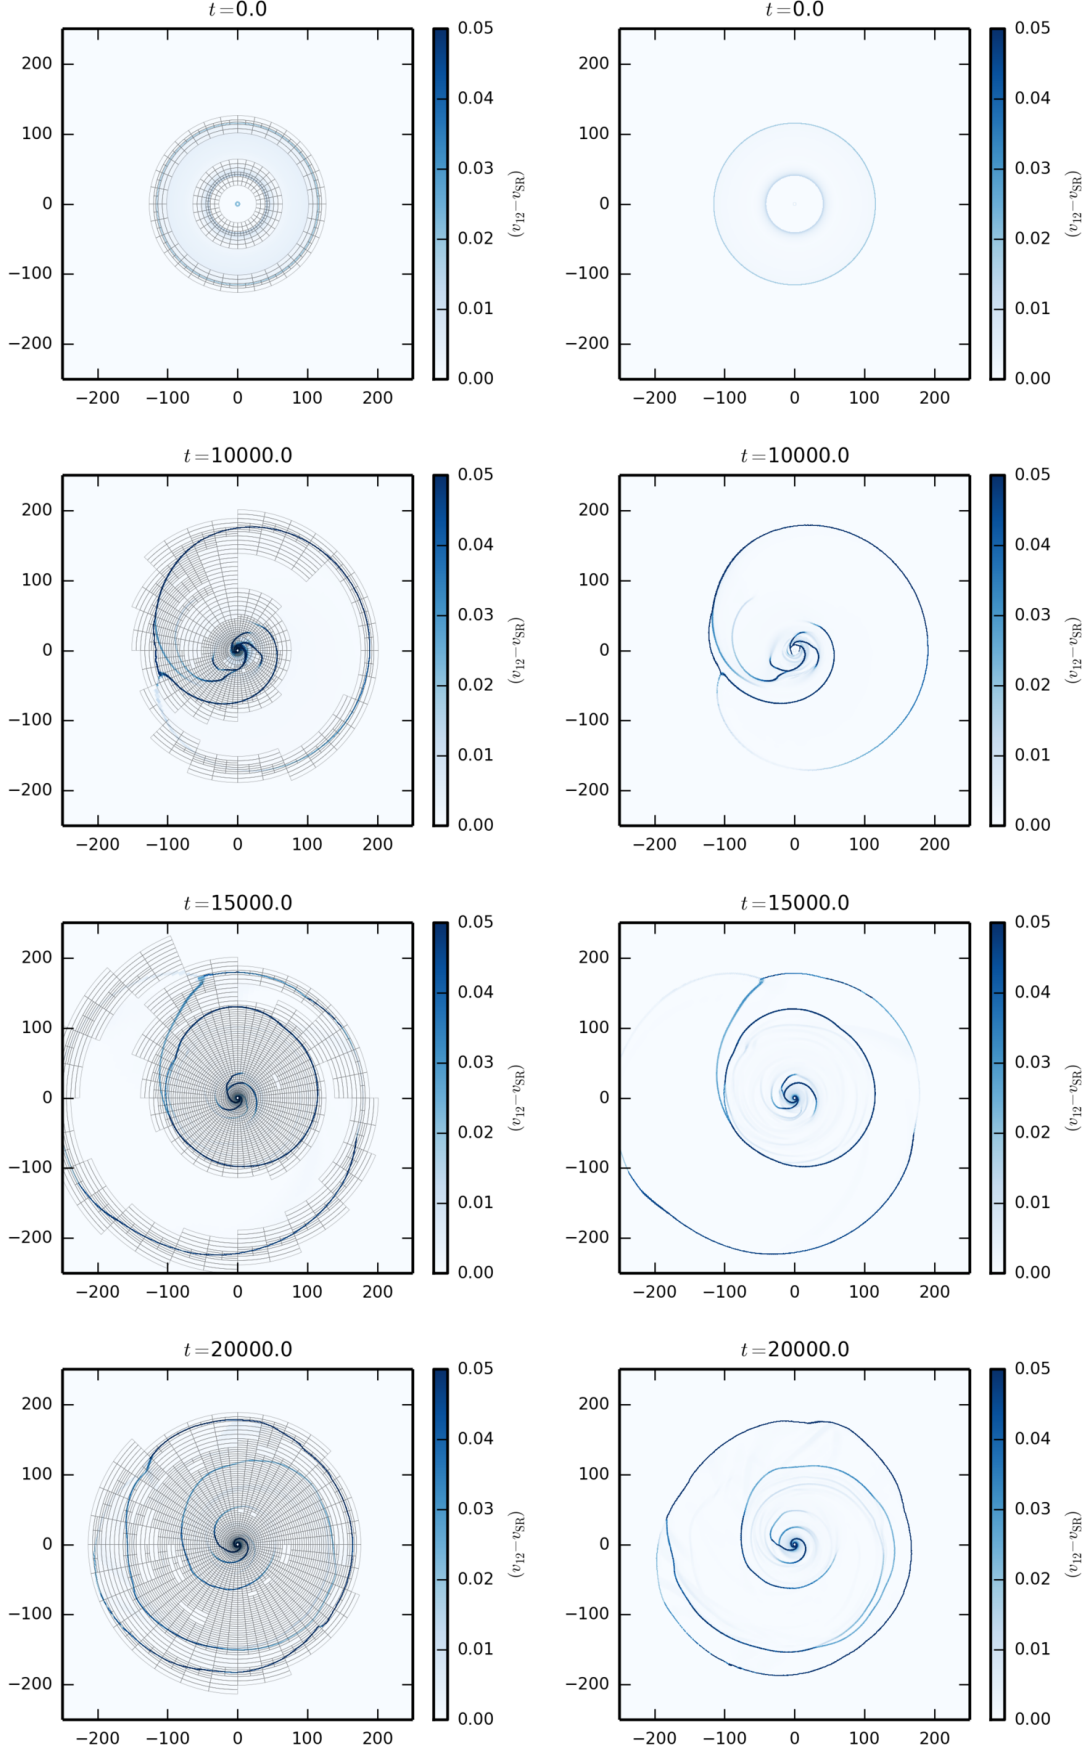


Fig. 9. Evolution of the shock structure in the 2D recoiling black hole simulations at times 0, 10,000, 15,000 and 20,000 M , with three AMR levels and tolerance $\varepsilon_t = 0.1$ (left columns) and a high resolution uniform grid (1024×512 , right columns). The colour indicates the difference between the relative velocity of the sides of the Riemann problem and the threshold velocity for producing a shock. For the AMR simulation, the blocks at levels higher than 1 are shown.




Fig. 10. Evolution of the internal energy of the disc normalised with respect to the initial internal energy of the disc and relative to the 2D simulations of a recoiling black hole. Different lines refer to simulations with uniform grid using three different resolutions: 256×128 (red solid line), 512×256 (green solid line), and 1024×512 (blue solid line). We also show the internal energy for AMR simulations using three levels with different tolerances of $\varepsilon_t = 0.1$ (blue dashed) and $\varepsilon_t = 0.005$ (blue dash-dotted).




Fig. 11. Total number of cells during the simulation of the 2D recoiling black hole with uniform grid using three different resolutions (256×128 : red, 512×256 : green, and 1024×512 : blue) and AMR with different tolerances of $\varepsilon_t = 0.1$ (two levels: green dashed, three levels: blue dashed) and $\varepsilon_t = 0.005$ (three levels: blue dash-dotted).

similar behaviour for the volume-integrated internal energy has been reported in [Megevand et al. \(2009\)](#). We also show in Fig. 10 the late-time behaviour of the volume-integrated internal energy, which changes slightly with grid resolution. However, the AMR cases are in very good agreement with the corresponding high-resolution run, demonstrating the successful capturing of the shock.

We also checked the convergence of the simulations with different resolutions for the uniform grid and AMR runs and obtained the expected convergence order in both cases (see Appendix A for more details).

As mentioned above, the value of the tolerance for switching on a refinement level has directly affects whether new cells are introduced where the equations are to be solved, thus translating into additional computational cost. Figure 11 shows the evolution of the total number of cells during each simulation for each of the different cases. The solid lines correspond to simulations

Table 1. CPU hours (CPUH) spent by the simulations of the 2D recoiling black hole at uniform resolutions, and fraction of that time spent by the equivalent AMR runs.

Grid size ($N_r \times N_\phi$)	CPU time uniform [CPUH]	Equiv. AMR time fraction [$\varepsilon_t = 0.1$]	Equiv. AMR time fraction [$\varepsilon_t = 0.005$]
256×128	55.0	—	—
512×256	443.1	0.65	0.70
1024×512	3,377.4	0.47	0.57

with uniform grids, while the dashed lines indicates AMR cases. Initially, $2^{17} \sim 131,000$ cells were used even when we used three AMR levels, which is similar to the number of cells for the simulation having uniform and medium resolution. When the simulations enter the accretion phase, however, the total number of cells rapidly increases because the spiral shock forms, triggering higher refinement and expanding within the disc. We note that the total number of cells is still smaller than or nearly half of the total number of cells in the corresponding high-resolution simulation (blue solid line), thus resulting in a direct reduction of the computational cost.

A comparison of the computational time for each of the 2D recoiling black hole simulations is shown in Table 1. It is remarkable that even the three-level AMR simulation could obtain results of an accuracy comparable to the high-resolution uniform run, but spent only slightly more than half the computational time used in the uniform run. In Sect. 4.3 we show that the agreement between this three-level AMR and its corresponding uniform run was also excellent for the general-relativistic radiative transfer calculation.

In summary, the AMR employed in our code has proven to be essential for physical scenarios such as the recoiling black hole, where the dynamics of the kicked accretion disc are very sensitive to the underlying numerical resolution. The AMR refinement strategy, triggered by the Löhner scheme, effectively captures the spiral shock structure developed in the accretion disc. Moreover, the simulations using AMR require only roughly half of the computational time of the corresponding uniform grid cases with highest resolution, making AMR a very useful tool for 3D simulations involving large-scale shocks.

3.4. Initial setup in 3D

In contrast to the previous section and to [Zanotti et al. \(2010\)](#), here we dropped the assumption that the disc is geometrically thin and evolved the dynamics in full 3D. The initial setup was now a geometrically thick torus with a constant angular momentum distribution as described in Sect. 2.3. The parameters of this torus and the black hole are the same as for the 2D case, so that the densities, pressures, and fluid velocities on the equatorial plane match those of the 2D simulation at $t = 0M$. Specifically, they are $a = 0.5$, $\ell = 8$, $r_{\text{in}} = 40M$, $r_{\text{out}} = 116M$, and $\hat{\gamma} = 4/3$.

The numerical domain extends over $r \in [1.85M, 400M]$ and $\phi \in [0, 2\pi]$. Taking advantage of the symmetry of the problem with respect to the equatorial plane, we considered only the upper half of the torus. The domain spanned the region $\theta \in [\pi/8, \pi/2]$. At the equatorial plane, symmetric boundary conditions were applied to all variables except for the vertical component of the velocity, for which an antisymmetric boundary condition was applied, all to account for a perfectly symmetric lower half of the torus. We note that during the simulation the fluid never reaches the outer boundaries of the domain. In this

case, the region outside of the torus was also filled with a tenuous atmosphere following the same prescription as employed in the 2D case.

We again quantified convergence and compared the performance of AMR to that of a high-resolution uniform grid simulation. To this end, we performed three simulations at uniform resolutions of $N_r = 128, 256$, and 512 , $N_\phi = \frac{1}{2}N_r$, and $N_\theta = \frac{1}{8}N_r$, to which we refer as low, medium, and high.

We performed three simulations using AMR with two and three levels, for which the base level had the same resolution as the low-resolution uniform run, and the highest level had the same resolution as the medium and high resolutions of the uniform cases, respectively. Owing to the higher computational cost of 3D simulations, this time we used AMR tolerances higher than in the 2D cases. Instead, for this setup we tested another feature of the implementation of AMR in the code, namely the possibility of specifying a different tolerance for triggering refinement at the various levels. Setting a higher tolerance for the highest levels results in a lower propensity of the code to refine towards those levels, which might decrease the computational cost.

More specifically, for the first two runs (2 and 3 levels), we set a tolerance of $\varepsilon_t = 0.1$ between every level. For the third run, we used three levels and set $\varepsilon_t = 0.1$ for refining to the second level and a less demanding 0.3 for refining to the third level.

In each case the system was evolved up to a time of $20,000 M$, corresponding to ~ 15 orbital periods.

3.5. Results in 3D

Figure 12 shows vertical (i.e. on the (x, z) plane) and horizontal (i.e. on the (x, y) plane) cuts of the density field (left half) and the shock structure (right half) of the torus at the final time of $t = 20,000 M$. We also show on the (x, y) and (y, z) planes the cuts that show the 3D location of the shocks produced by the recoiling black hole as it interacts with the accreting torus.

Here the kick velocity also breaks the symmetry of the initial density profile and leads to an accumulation of gas in a small region of the disc. Similarly to the 2D case, the compression of gas due to the kick velocity eventually evolves into a spiral shock that is completely visible around $t = 10,000 M$. As can be appreciated in the vertical slices of Fig. 12, now the spiral shock also extends in the vertical direction. However, in contrast to the 2D case, accretion onto the central black hole does not start until $t \approx 19,500 M$. The reason is that the fluid is no longer confined to the equatorial plane when the interaction with the black hole produces a compression wave that, together with the centrifugal force, allows the fluid to expand in the vertical direction as well.

In analogy with the 2D case, Fig. 13 shows the volume-integrated internal energy of the torus, normalised to its initial value. By comparing it with Fig. 10, we can appreciate some differences between the dynamics in 2D and 3D. While in 2D the internal energy falls at around $t = 10,000 M$ as a result of the expansion of the disc and rises again to a nearly constant value when the gas is heated by the shock, in 3D this is no longer possible because the fluid cools down when it expands in the vertical direction as a result of the interaction with the shock.

Even though the evolution of the internal energy of the simulation at the lowest resolution differs significantly from the results of the other simulations, it is still qualitatively similar. In Appendix A we show that in uniform grids as well as in AMR, the solution of the equations converges at the expected order.

Table 2 is the equivalent of Table 1 for the 3D simulations and shows the CPU time spent by simulations performed at uni-




Fig. 13. Evolution of volume-integrated (normalised) internal energy of the torus for the 3D recoiling black hole cases, with uniform grid (solid) using the two highest resolutions ($256 \times 128 \times 32$: green, and $512 \times 256 \times 64$: blue) and AMR (dashed) with two (green) and three (blue) levels.

Table 2. CPU hours (CPUH) spent by the simulations of the 3D recoiling black hole at uniform resolutions, and fraction of that time spent by the equivalent AMR runs.

Grid size $[N_r \times N_\phi \times N_\theta]$	CPU time uniform [CPUH]	Equiv. AMR time fraction $[\varepsilon_t = 0.1, 0.1]$	Equiv. AMR time fraction $[\varepsilon_t = 0.1, 0.3]$
$128 \times 64 \times 16$	667.1	—	—
$256 \times 128 \times 32$	8,557.2	0.15	—
$512 \times 256 \times 64$	93,144.8	0.14	0.14

form resolutions and the fraction of that time spent by simulations using AMR. Two three-level AMR runs were performed that differ only in the refinement threshold on the highest level: $\varepsilon_t = 0.1$ with respect to $\varepsilon_t = 0.3$. While the results were practically indistinguishable from one another, only a very small saving in computational time was achieved, namely 13,550.0 versus 13,416.0 CPU-hours, which corresponds to a difference of 0.5 %. The AMR simulations obtained an accuracy comparable to the equivalent uniform-resolution simulations, but with a saving in CPU time of $\sim 85\%$. In other words, we obtained numerically equivalent results using slightly more than 1/10 of the resources. Even though the three-level AMR run is closer to the medium-resolution run than to the high-resolution run, both AMR simulations remain close to their equivalent uniform runs. As a final remark, we emphasise again that as was shown for the 2D case, an AMR simulation can become as close as desired to its equivalent high-resolution uniform run by reducing the tolerance ε_t .

4. Ray-tracing and radiation transfer of solutions

To accurately compute the electromagnetic emissions from our simulations, it is necessary to perform ray-tracing calculations coupled with general-relativistic radiation transfer calculations (e.g. Fuerst & Wu 2004; Vincent et al. 2011; Younsi et al. 2012; Younsi & Wu 2015; Dexter 2016; Pu et al. 2016). These calculations were performed in post-processing and therefore the effect of radiation forces coupled with the hydrodynamic evolution of the material were not included. We also employed the so-called fast-light approximation, where the dynamical timescale of the simulation is taken to be much longer than the light-crossing time, and so the finite travel time of photons and their relative

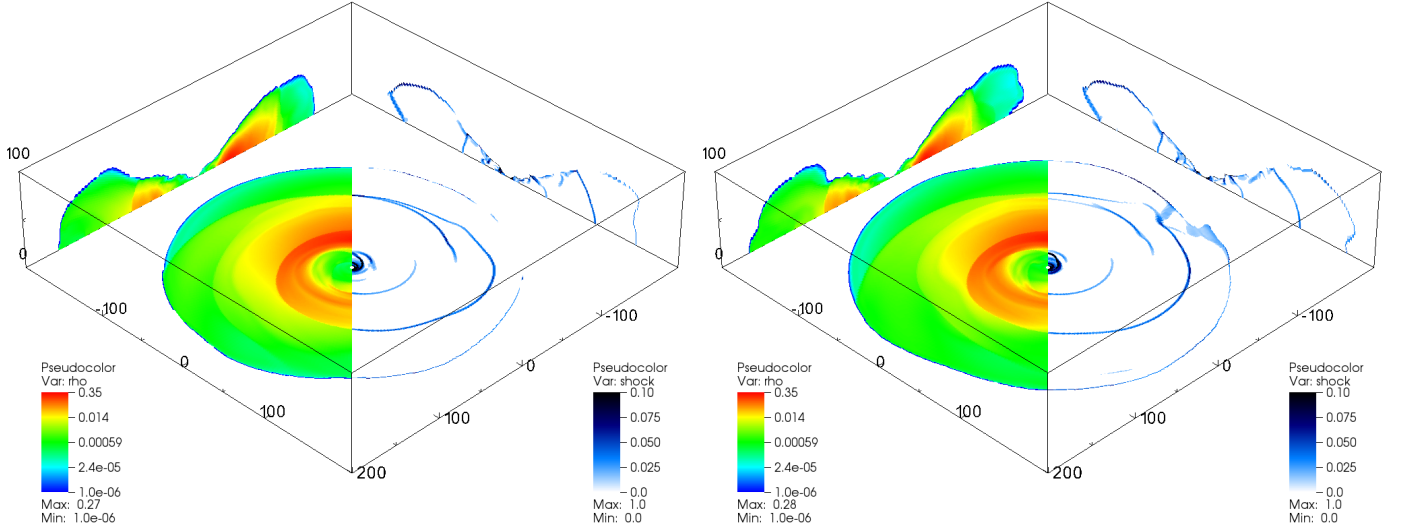


Fig. 12. Snapshots of the logarithmic density and the shock structure of the 3D recoiling black hole at the final time, $t = 20,000 M$, for the uniform run with resolution $512 \times 256 \times 64$ (left) and an equivalent three-level AMR run with high tolerance ($\epsilon_i = 0.1, 0.3$, right). The floor of the box shows a cut through the equatorial plane and the walls show perpendicular slices crossing the origin.

arrival time delays may be neglected. Such an approximation is acceptable for the large-scale recoiling black hole simulations considered in this paper.

Electromagnetic radiation follows null geodesics of the space-time, thus we calculated the geodesics through direct numerical integration of the geodesic equations of motion. The geodesics were solved using an adaptive fourth-order Runge-Kutta scheme, integrating backwards in time from an observer at $10^3 r_g$ from the black hole (where space-time is practically Euclidean) and assuming that all rays arrive perpendicular to the observer's image plane.

After calculating the geodesic for each ray, we then solved the radiation transport equation. We employed the ray-tracing and radiation transport scheme described in [Younsi et al. \(2012\)](#). In covariant form, the general-relativistic radiation transport equation (in the absence of scattering) may be written as

$$\frac{d\mathcal{I}}{d\lambda} = -k_\mu u^\mu|_\lambda \left(-\alpha_{\nu,0} \mathcal{I} + \frac{j_{\nu,0}}{\nu^3} \right), \quad (15)$$

where the Lorentz-invariant intensity $\mathcal{I} \equiv I_\nu/\nu^3$, ν is the frequency of radiation, I_ν is the specific intensity, and $\alpha_{\nu,0}$ and $j_{\nu,0}$ are the specific absorption and emission coefficient, evaluated at frequency ν and in the local fluid rest frame (hence denoted by the subscript 0). Here k_μ is the photon four-momentum, u^μ is the four-velocity of the emitting medium, and λ is the affine parameter. This equation may be rewritten in terms of the optical depth of the medium as

$$\frac{d\mathcal{I}}{d\tau_\nu} = -\mathcal{I} + \frac{\eta}{\chi}, \quad (16)$$

where τ_ν , the optical depth evaluated at frequency ν , is calculated as

$$\tau_\nu(\lambda) = - \int_{\lambda_0}^{\lambda} d\lambda' \alpha_{\nu,0}(\lambda') k_\mu u^\mu|_{\lambda'}, \quad (17)$$

and the invariant emission coefficient η and invariant absorption coefficient χ are defined as $\eta \equiv j_\nu/\nu^2$ and $\chi \equiv \nu\alpha_\nu$.

The radiative-transfer equation (16) may itself be reduced to two differential equations (see [Younsi et al. 2012](#), for details), yielding

$$\frac{d\tau_\nu}{d\lambda} = \gamma^{-1} \alpha_{\nu,0}, \quad (18)$$

$$\frac{d\mathcal{I}}{d\lambda} = \gamma^{-1} \frac{j_{\nu,0}}{\nu^3} \exp(-\tau_\nu), \quad (19)$$

where the relative energy shift, γ (not to be confused with the determinant of the three-metric), between the radiation emitted from material orbiting the black hole with four-velocity u^α and the radiation received by a distant observer is given by

$$\gamma^{-1} \equiv \frac{\nu_0}{\nu} = \frac{k_\alpha u^\alpha|_0}{k_\beta u^\beta|_{\text{obs}}}. \quad (20)$$

The subscript obs denotes the reference frame of a distant observer. Given that the background metric is stationary, the geodesic equations of motion are time-symmetric. The fast-light approximation was also adopted, therefore the fluid at each observer time slice is stationary, and Eqs. (18)-(19) are also time-symmetric. We consequently set both the initial intensity \mathcal{I} and initial opacity τ_{nu} to zero, directly integrating Eqs. (18)-(19) together with the geodesic equations of motion, backwards in time.

We illustrate the features of this approach. Firstly, it avoids the process of having to integrate the geodesics backwards in time, store the geodesics in memory, and then integrate the radiative transfer equations forward in time towards the observer. Secondly, it offers the option of specifying a threshold optical depth (typically on the order of unity) when encountering optically thick media, enabling the geodesic integration to be terminated when this optical depth threshold is exceeded. Consequently, this approach saves significant computational expense and time.

4.1. Thermodynamic quantities

When we calculate the radiation transport of simulation data, we must specify the emission and absorption coefficients for all relevant radiative processes. These coefficients must be calculated in physical units, whereas the simulation data are output in geometrised units. Length and times are easily converted into cgs units through re-introducing the mass, M , of the black hole, which hereafter we take to be $M = 10^8 M_\odot$. However, the emission and absorption coefficients also depend on the density and temperature in cgs units. Following [Schnittman et al. \(2013\)](#), and using the fact that our EOS is ideal, the conversion between geometrised (geo) and cgs units for the fluid temperature is given by

$$T_{\text{cgs}} = \left(\frac{P_{\text{geo}}}{\rho_{\text{geo}}} \right) \frac{\mu m_p}{k_B} c^2, \quad (21)$$

where μ is the mean molecular weight of electrons and ions, m_p is the proton rest mass, and k_B is the Boltzmann constant. In analogy to [Anderson et al. \(2010\)](#) and [Zanotti et al. \(2010\)](#), we scaled the initial rest mass density at the centre of the torus to be $\rho_c = 1.38 \times 10^{-10} \text{ g cm}^{-3}$. The mean molecular weight was determined from

$$\frac{1}{\mu} \equiv \frac{1}{\mu_e} + \frac{1}{\mu_i}, \quad (22)$$

where the effective molecular weights of electrons and ions are given by

$$\mu_e \equiv \frac{2}{1+X}, \quad \mu_i \equiv \frac{4}{1+3X}. \quad (23)$$

Consequently, the mean molecular weight is given by

$$\mu = \frac{4}{3+5X}, \quad (24)$$

where X , the relative abundance of hydrogen, is set to $3/4$ in all our calculations, giving $\mu_e = 8/7$, $\mu_i = 16/13$, and $\mu = 16/27$. With these definitions the electron and ion number densities are given by

$$n_e = \frac{\rho_{\text{cgs}}}{\mu_e m_p}, \quad n_i = \frac{\rho_{\text{cgs}}}{\mu_i m_p}. \quad (25)$$

4.2. Radiative parameters

To calculate the electromagnetic emission from the recoiling black-hole simulations, we assumed emission primarily in the form of thermal bremsstrahlung from electron-ion and electron-electron interactions. Owing to the relativistic equation of state used in these simulations and because the temperature can range between $\sim 10^7 - 10^{11} \text{ K}$, it is necessary to employ a relativistic Maxwellian distribution for the population of thermal electrons. Following [Stepney & Guilbert \(1983\)](#) and [Narayan & Yi \(1995\)](#), the total thermal bremsstrahlung cooling rate (see also [Straub et al. 2012](#)) may be written as

$$q_{\text{br}}^- = q_{\text{ei}}^- + q_{\text{ee}}^-, \quad (26)$$

where q_{ei}^- and q_{ee}^- are the electron-ion and electron-electron cooling terms, respectively. The electron-ion cooling rate is given by

$$q_{\text{ei}}^- = 3.013 \times 10^{25} \rho_{\text{cgs}}^2 F_{\text{ei}}(\Theta_e) \text{ erg cm}^{-3} \text{ s}^{-1}, \quad (27)$$

where the dimensionless electron temperature (Θ_e) is defined as

$$\Theta_e \equiv \frac{k_B T_e}{m_e c^2}, \quad (28)$$

and where T_e is the electron temperature and m_e the electron rest mass. The function $F_{\text{ei}}(\Theta_e)$ is given by

$$F_{\text{ei}}(\Theta_e) = \begin{cases} 4 \left(\frac{2\Theta_e}{\pi^3} \right)^{1/2} (1 + 1.781\Theta_e^{1.34}), & \Theta_e < 1, \\ \frac{9\Theta_e}{2\pi} [1.5 + \ln(1.123\Theta_e + 0.48)], & \Theta_e > 1, \end{cases} \quad (29)$$

where F_{ei} is continuous across $\Theta_e = 1$. The electron-electron bremsstrahlung cooling term is given by

$$q_{\text{ee}}^- = \begin{cases} C_1 \rho_{\text{cgs}}^2 \Theta_e^{3/2} (1 + 1.1\Theta_e + \Theta_e^2 - 1.25\Theta_e^{5/2}), & \Theta_e < 1, \\ C_2 \rho_{\text{cgs}}^2 \Theta_e (\ln 1.123\Theta_e + 1.28), & \Theta_e > 1, \end{cases} \quad (30)$$

where $C_1 = 7.028 \times 10^{25}$ and $C_2 = 9.334 \times 10^{25}$, and q_{ee}^- has units of $\text{erg cm}^{-3} \text{ s}^{-1}$. With the thermal bremsstrahlung cooling rates in hand, we may now write the total emissivity in the fluid rest frame as

$$j_{\nu, \text{br}} = \frac{1}{4\pi\nu} q_{\text{br}}^- x e^{-x} \bar{g}(x), \quad (31)$$

where

$$x \equiv \frac{h_p \nu}{k_B T_e}, \quad (32)$$

and h_p is the Planck constant. The factor of $1/(4\pi)$ specifies isotropic emission in the fluid rest frame, and the mean Gaunt factor, $\bar{g}(x)$, is given by

$$\bar{g}(x) \equiv \begin{cases} \frac{\sqrt{3}}{\pi} \ln(2.246x^{-1}), & x < 1, \\ \left(\frac{3}{\pi} x^{-1} \right)^{1/2}, & x > 1. \end{cases} \quad (33)$$

In all calculations reported here, we assumed that the ionic contribution comes exclusively from protons. The simulation data provide only the equilibrium temperature of electrons and protons, and not of individual species, therefore we take $T_e = T_{\text{cgs}}$, assuming that the local electron and proton temperatures do not differ significantly from the local equilibrium temperature.

In 3D models where we considered the opacity of the emitting medium, we assumed a modified Kramer opacity law as employed in [Schnittman et al. \(2006\)](#) and [Anderson et al. \(2010\)](#), where

$$\alpha_{0,\nu} = 5 \times 10^{24} \rho_{\text{cgs}}^2 T^{-7/2} \left(\frac{1 - e^{-x}}{x^3} \right) \text{ cm}^{-1}. \quad (34)$$

This opacity adds thermal radiation for optically thick regions, whilst the optically thin regions radiate bremsstrahlung (see [Anderson et al. 2010](#)).




Fig. 14. Ray-tracing and radiative transfer calculation of 2D recoiling black-hole simulation. Left column: as viewed at an observer inclination angle of $\theta_{\text{obs}} = 0.1^\circ$. Right column: viewed at an observer inclination angle of $\theta_{\text{obs}} = 60^\circ$. From top row to bottom row, as viewed at $t = 0, 10,000, 15,000$, and $20,000 M$. The colour scale is logarithmic in the total intensity, I , from each pixel (arbitrary units).

4.3. Recoiling black hole in 2D

The dynamics of the 2D recoiling black hole is ultimately that of a planar flow in the equatorial ($\theta = \pi/2$) plane of the black hole. When we ray-trace these simulations, we need only calculate the

intersection point of the ray with the equatorial plane, determining the emitted spectrum at that particular pixel of the image. Since the ray does not traverse the emitting medium, the emission is optically thick and planar, thus the calculated results are scale-free and do not depend on the mass of the black hole. To




Fig. 15. Top and middle panels: normalised total flux light curves of the emission from the 2D recoiling black hole simulation for $\theta_{\text{obs}} = 0.1^\circ$ and $\theta_{\text{obs}} = 60^\circ$, respectively. Solid lines are for the simulation run with three AMR levels and a tolerance of $\varepsilon_t = 0.005$. Dashed lines are for the uniform grid simulation run of equivalent resolution. Bottom: flux difference between the 2D AMR and uniform run light-curves for $\theta_{\text{obs}} = 0.1^\circ$ (solid) and $\theta_{\text{obs}} = 60^\circ$ (dashed).

generate each image, we ray-traced a grid of $2,500 \times 2,500$ photons, sampling 200 uniformly logarithmically spaced frequency bins between 10^5 Hz and 10^{25} Hz. Each pixel of the calculated images represents the total frequency-integrated emission.

We also present calculations of light curves for different inclination parameters, where the total integrated intensity over all frequencies and over every pixel in an image (i.e. flux) corresponds to that point in time on the light curve. We recall that the intensity is the energy received per unit time, and we also refer to it as the “flux”, which should not be confused with the “fluxes” introduced in Eq. (7), however. Since we stored the entire spectrum for each pixel, we can also calculate the image and light curve at specific observer frequencies, which is of practical interest when comparing images from radiation transport calculations of GRMHD simulation data with observations of Sagittarius A* (Sgr A*) at 1.3 mm, for instance.

Figure 14 presents radiation image calculations of the 2D recoiling black hole simulation. For an inclination angle of $\theta_{\text{obs}} = 0.1^\circ$, the structure of the flow is similar to the renderings in Fig. 8. One obvious difference is the absence of emission from the innermost region in the vicinity of the black hole event horizon. When instead $\theta_{\text{obs}} = 60^\circ$, the image of the 2D recoiling black hole is warped and has a smaller projected surface area, with the approaching side of the flow being Doppler-boosted and projected along the line of sight, and conversely for the receding side.

In Fig. 15 we present light-curve calculations of the AMR and uniform grid runs of the same 2D recoiling black hole simulation. As in Fig. 14, we considered two observer inclination angles and calculated the light curves from the two sets of simulation data. The AMR and uniform grid runs are in excellent agreement, indicating that the AMR simulation captures both the qualitative and quantitative aspects of the dynamics and thermodynamics well, which is reflected in the light curves. The bottom panel reveals that while the differences between the two runs vary, they always remain below the 1% level. For an observer at $\theta_{\text{obs}} = 0.1^\circ$, the oscillations are almost exclusively due to the growth and propagation of the large spiral shock, since Doppler and aberrational effects are essentially uniform across the entire image at such low inclinations. For an inclination $\theta_{\text{obs}} = 60^\circ$,




Fig. 17. As in Fig. 15, but now considering the 3D recoiling black hole simulation. The higher AMR level tolerance for the 3D run ($\varepsilon_t = 0.1, 0.3$) results in a less perfect agreement with the uniform-resolution run than for 2D ($\varepsilon_t = 0.005$, Fig. 15), as is discussed in Sect. 3.4.

the flux is lower (since the projected surface area is smaller) and more oscillations of the light-curve can be seen. These are due to the additional presence of non-uniform Doppler boosting of the emission from the spiral-shocked material as it approaches the observer (peaks) and as it moves away from the observer (troughs).

4.4. Recoiling black hole in 3D

In analogy with Anderson et al. (2010) and Zanotti et al. (2010), we set the initial rest-mass density at the centre of the torus to be $\rho_c = 1.38 \times 10^{-10} \text{ g cm}^{-3}$. The image parameters are identical to the 2D case. Figure 16 presents radiation image calculations of the 3D recoiling black hole simulation for the AMR run. For $\theta_{\text{obs}} = 0.1^\circ$ the panels in Figs. 14 and 16 appear similar, but there are several differences.

The first difference is the near-absence of accretion at early times in 3D, occurring (much more slowly) at later times than in the 2D case. The second difference is the lensed emission from the torus, which manifests itself as the inner ring of emission and is most distinct at $t = 0$. This is due to rays that traverse the entire 3D emitting medium multiple times before reaching the observer. For an observer at $\theta_{\text{obs}} = 60^\circ$, the front limb of the torus obscures the central region of the black hole and opacity effects dominate, giving rise to much stronger emission from this region than in either the $\theta_{\text{obs}} = 0.1^\circ$ case in 3D or all viewing angles in the 2D case (where self-obscuration is absent). A third major difference is that, particularly at late times, shocked regions and the emission from near the vicinity of the event horizon appear more optically thick because the flow is 3D and absorptive as well as emissive.

In Fig. 17 we present light-curve calculations of the AMR and uniform grid runs of the same 3D recoiling black hole simulation. Unlike the 2D case that we investigated previously, the tolerance in this simulation was higher, i.e. $\varepsilon_t = 0.1$ (see discussion in Sect. 3.4). As expected, the higher tolerance causes larger differences between the uniform and AMR runs, the maximum difference being $\sim 9.8\%$. However, the light curves still retain the same morphological profiles and relative properties, with the AMR runs slightly overestimating the flux relative to the uniform run. These differences are acceptable and do not change the physical conclusions drawn from these calculations. Moreover, considering that the difference in runtime between the uniform




Fig. 16. Ray-tracing and radiative-transfer calculation of the 3D recoiling black hole simulation. Same panel descriptions as Fig. 14, but now the colour scale is logarithmic in I/I_{max} , i.e. normalised to the peak intensity.

and AMR simulations was a factor of ~ 7 , this represents a significant speedup, allowing us to employ still higher resolutions.

5. Conclusions

We have discussed results from a new 3D general-relativistic hydrodynamics code with grid-based AMR capabilities, the motivation for which arose mainly from our own continued efforts in

augmenting the wealth of community codes available for astrophysical research.

The code was tested in the general-relativistic regime by evolving a number of stationary and non-stationary flows onto black hole space-times, including the spherical (Michel) accretion onto a Schwarzschild black hole and stationary tori with a constant angular momentum in a rotating black hole space-time using Boyer-Lindquist and Kerr-Schild coordinates. We further demonstrated that the code can be properly employed in the consideration of other scientific applications.

A particularly critical test performed has involved the evolution in 2D and in 3D of a black hole recoiling into a circumbinary accretion disc, where both the nonlinearity of the dynamics and the development of strong large-scale shocks have been tested, making use of the capabilities of AMR. In particular, we have shown that AMR is essential for recoiling black hole simulations because the dynamics of the kicked accretion disc are very sensitive to the numerical resolution, and AMR has proven effective in capturing and resolving the spiral shock structure that develops in the accretion disc. AMR has also been shown to be very economical, requiring only half of the computational grid and time compared to the high-resolution case without AMR and still yielding virtually unchanged results.

Our relativistic hydrodynamics calculations have also been coupled to a consistent treatment of the general-relativistic radiation-transport equation to compute the electromagnetic emissions from the underlying dynamics of the flow. The radiative-emission calculations were performed in post-processing and combined with ray-tracing techniques to obtain a somewhat realistic representation of the electromagnetic emission from this process for the first time.

In summary, the work presented here lays the ground for the development of a generic computational infrastructure to accurately and self-consistently calculate accretion flows onto compact objects, either black holes or neutron stars, and to compute with an increased degree of precision the associated electromagnetic emission from these scenarios. This could have a direct effect on collaborative efforts such as the Event Horizon Telescope Collaboration¹ (Doeleman et al. 2009) or the Black Hole Camera project² (Goddi et al. 2016). Work is already ongoing to include the effects of magnetic fields in the ideal-magnetohydrodynamics limit and will be presented in a forthcoming publication.

Acknowledgements. It is a pleasure to thank M. De Laurentis and C. Fromm for discussion and comments. This research is supported by the ERC Synergy Grant “BlackHoleCam – Imaging the Event Horizon of Black Holes” (Grant 610058). ZY is supported by an Alexander von Humboldt Fellowship. HO gratefully acknowledges the support from a CONACYT-DAAD scholarship. The simulations were performed on LOEWE at the CSC-Frankfurt.

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, *Phys. Rev. Lett.*, 116, 061102
- Alic, D., Moesta, P., Rezzolla, L., Zanotti, O., & Jaramillo, J. L. 2012, *Astrophys. J.*, 754, 36
- Amaro-Seoane, P., Aoudia, S., Babak, S., et al. 2012, *Class. and Quantum Grav.*, 29, 124016
- Anderson, M., Lehner, L., Megevand, M., & Neilsen, D. 2010, *Phys. Rev. D*, 81, 044004
- Anninos, P., Fragile, P. C., & Salmonson, J. D. 2005, *Astrophys. J.*, 635, 723
- Antón, L., Zanotti, O., Miralles, J. A., et al. 2006, *Astrophys. J.*, 637, 296
- Armitage, P. J. & Natarajan, P. 2002, *Astrophys. J.*, 567, L9
- Baiotti, L., Hawke, I., Montero, P. J., et al. 2005, *Phys. Rev. D*, 71, 024035
- Baker, J. G., Boggs, W. D., Centrella, J., et al. 2007, *Astrophys. J.*, 668, 1140
- Banyuls, F., Font, J. A., Ibáñez, J. M., Martí, J. M., & Miralles, J. A. 1997, *Astrophys. J.*, 476, 221
- Barausse, E., Morozova, V., & Rezzolla, L. 2012, *Astrophys. J.*, 758, 63
- Bode, T., Bogdanović, T., Haas, R., et al. 2012, *Astrophys. J.*, 744, 45
- Bondi, H. 1952, *Mon. Not. R. Astron. Soc.*, 112, 195
- Calder, A. C., Fryxell, B., Plewa, T., et al. 2002, *Astrophys. J., Supp.*, 143, 201
- Campanelli, M., Lousto, C. O., Zlochower, Y., & Merritt, D. 2007, *Physical Review Letters*, 98, 231102
- Colella, P. & Woodward, P. R. 1984, *Journal of Computational Physics*, 54, 174
- Corrales, L. R., Haiman, Z., & MacFadyen, A. 2010, *Mon. Not. R. Astron. Soc.*, 404, 947
- Cunningham, A. J., Frank, A., Varnière, P., Mitran, S., & Jones, T. W. 2009, *Astrophys. J., Supp.*, 182, 519
- De Villiers, J.-P. & Hawley, J. F. 2003, *Astrophys. J.*, 589, 458
- Del Zanna, L., Zanotti, O., Bucciantini, N., & Londrillo, P. 2007, *Astron. Astrophys.*, 473, 11
- Dexter, J. 2016, *MNRAS*, 462, 115
- Dibi, S., Drappeau, S., Fragile, P. C., Markoff, S., & Dexter, J. 2012, *Mon. Not. R. Astron. Soc.*, 426, 1928
- Dionysopoulou, K., Alic, D., Palenzuela, C., Rezzolla, L., & Giacomazzo, B. 2013, *Phys. Rev. D*, 88, 044020
- Doeleman, S., Agol, E., Backer, D., et al. 2009, in *Astronomy*, Vol. 2010, astro2010: The Astronomy and Astrophysics Decadal Survey
- Duez, M. D., Liu, Y. T., Shapiro, S. L., & Stephens, B. C. 2005, *Phys. Rev. D*, 72, 024028
- Etienne, Z. B., Paschalidis, V., Haas, R., Mösta, P., & Shapiro, S. L. 2015, *Classical and Quantum Gravity*, 32, 175009
- Farris, B. D., Gold, R., Paschalidis, V., Etienne, Z. B., & Shapiro, S. L. 2012, *Phys. Rev. Lett.*, 109, 221102
- Farris, B. D., Liu, Y. T., & Shapiro, S. L. 2010, *Phys. Rev. D*, 81, 084008
- Farris, B. D., Liu, Y. T., & Shapiro, S. L. 2011, *Phys. Rev. D*, D84, 024024
- Fishbone, L. G. & Moncrief, V. 1976, *Astrophys. J.*, 207, 962
- Font, J. A. 2003, *Living Rev. Relativ.*, 6, 4
- Font, J. A. & Daigne, F. 2002, *Astrophys. J.*, 581, L23
- Foucart, F., Chandra, M., Gammie, C. F., & Quataert, E. 2016, *Mon. Not. R. Astron. Soc.*, 456, 1332
- Fragile, P. C., Olejar, A., & Anninos, P. 2014, *Astrophys. J.*, 796, 22
- Fuerst, S. V. & Wu, K. 2004, *Astron. Astrophys.*, 424, 733
- Galeazzi, F., Kastaun, W., Rezzolla, L., & Font, J. A. 2013, *Phys. Rev. D*, 88, 064009
- Gammie, C. F., McKinney, J. C., & Tóth, G. 2003, *Astrophys. J.*, 589, 458
- Giacomazzo, B., Baker, J. G., Miller, M. C., Reynolds, C. S., & van Meter, J. R. 2012, *Astrophys. J.*, 752, L15
- Giacomazzo, B. & Rezzolla, L. 2007, *Class. Quantum Grav.*, 24, 235
- Goddi, C., Falke, H., Kramer, M., Rezzolla, L., & et al. 2016, *International Journal of Modern Physics D [arXiv:1606.08879]*, submitted
- Gold, R., Paschalidis, V., Etienne, Z. B., Shapiro, S. L., & Pfeiffer, H. P. 2014a, *Phys. Rev. D*, 89, 064060
- Gold, R., Paschalidis, V., Ruiz, M., et al. 2014b, *Phys. Rev. D*, 90, 104030
- Gonzalez, J. A., Sperhake, U., Bruegmann, B., Hannam, M., & Husa, S. 2007, *Phys. Rev. Lett.*, 98, 091101
- Graham, M. J., Djorgovski, S. G., Stern, D., et al. 2015, *Nature*, 518, 74
- Haehnelt, M. G. 1994, *Mon. Not. R. Astron. Soc.*, 269, 199
- Hamlin, N. D. & Newman, W. I. 2013, *Phys. Rev. E*, 87, 043101
- Harten, A., Lax, P. D., & van Leer, B. 1983, *SIAM Rev.*, 25, 35
- Hawley, J. F., Smarr, L. L., & Wilson, J. R. 1984, *Astrophys. J.*, 277, 296
- Keppens, R., Meliani, Z., van Marle, A. J., et al. 2012, *Journal of Computational Physics*, 231, 718
- Komossa, S. 2012, *Advances in Astronomy*, 2012, 364973
- Komossa, S., Burwitz, V., Hasinger, G., et al. 2003, *Astrophys. J. Lett.*, 582, L15
- Koppitz, M. et al. 2007, *Phys. Rev. Lett.*, 99, 041102
- Koren, B. 1993, *Numerical methods for advection–diffusion problems*, ed. C. B. Vreugdenhil & B. Koren., Notes on numerical fluid mechanics, v. 45 (Braunschweig: Vieweg)
- Kormendy, J. & Ho, L. C. 2013, *Annual Review of Astronomy and Astrophysics*, 51, 511
- Kozłowski, M., Jaroszynski, M., & Abramowicz, M. A. 1978, *Astron. and Astrophys.*, 63, 209
- Kudoh, S. K. D. M. K. S. T. 2000, *Astrophys. J.*, 536, 668
- Lippai, Z., Frei, Z., & Haiman, Z. 2008, *Astrophys. J.*, 676, L5
- Löhner, R. 1987, *Computer Methods in Applied Mechanics and Engineering*, 61, 323
- Martí, J. M. & Müller, E. 2015, *Living Reviews in Computational Astrophysics*, 1
- McKinney, J. C., Tchekhovskoy, A., Sadowski, A., & Narayan, R. 2014, *Mon. Not. R. Astron. Soc.*, 441, 3177
- Megevand, M., Anderson, M., Frank, J., et al. 2009, *Phys. Rev. D*, 80, 024012
- Michel, F. C. 1972, *Astrophys. Spa. Sci.*, 15, 153

Mignone, A., Zanni, C., Tzeferacos, P., et al. 2012, *Astrophys. J. Suppl. Ser.*, 198, 7

Milosavljeć, M. & Phinney, E. S. 2005, *Astrophys. J.*, 622, L93

Milosavljević, M. & Merritt, D. 2001, *Astrophys. J.*, 563, 34

Mizuno, Y., Nishikawa, K.-I., Koide, S., Hardee, P., & Fishman, G. J. 2006, *ArXiv Astrophysics e-prints [astro-ph/0609004]*

Moesta, P., Alic, D., Rezzolla, L., Zanotti, O., & Palenzuela, C. 2012, *Astrophys. J. Lett.*, 749, L32

Mösta, P., Palenzuela, C., Rezzolla, L., et al. 2010a, *Phys. Rev. D*, 81, 064017

Mösta, P., Palenzuela, C., Rezzolla, L., et al. 2010b, *Phys. Rev. D*, 81, 064017

Narayan, R. & Yi, I. 1995, *Astrophys. J.*, 452, 710

Noble, S. C., Gammie, C. F., McKinney, J. C., & Del Zanna, L. 2006, *Astrophys. J.*, 641, 626

Noble, S. C., Mundim, B. C., Nakano, H., et al. 2012, *Astrophys. J.*, 755, 51

O’Neill, S. M., Miller, M. C., Bogdanović, T., Reynolds, C. S., & Schnittman, J. D. 2009, *Astrophys. J.*, 700, 859

Palenzuela, C., Anderson, M., Lehner, L., Liebling, S. L., & Neilsen, D. 2009, *Phys. Rev. Lett.*, 103, 081101

Palenzuela, C., Lehner, L., & Liebling, S. L. 2010, *Science*, 329, 927

Palenzuela, C., Lehner, L., & Yoshida, S. 2010, *Phys. Rev. D*, 81, 084007

Ponce, M., Faber, J. A., & Lombardi, J. C. 2012, *Astrophys. J.*, 745, 71

Porth, O., Xia, C., Hendrix, T., Moschou, S. P., & Keppens, R. 2014, *Astrophys. J. Suppl.*, 214, 4

Pu, H.-Y., Yun, K., Younsi, Z., & Yoon, S.-J. 2016, *Astrophys. J.*, 820, 105

Radice, D. & Rezzolla, L. 2012, *Astron. Astrophys.*, 547, A26

Radice, D., Rezzolla, L., & Galeazzi, F. 2014, *Mon. Not. R. Astron. Soc. L.*, 437, L46

Reisswig, C., Husa, S., Rezzolla, L., et al. 2009, *Phys. Rev. D*, 80, 124026

Rezzolla, L. 2009, *Class. Quantum Grav.*, 26, 094023

Rezzolla, L. & Zanotti, O. 2002, *Phys. Rev. Lett.*, 89, 114501

Rezzolla, L. & Zanotti, O. 2013, *Relativistic Hydrodynamics* (Oxford, UK: Oxford University Press)

Rezzolla, L., Zanotti, O., & Pons, J. A. 2003, *Journ. of Fluid Mech.*, 479, 199

Rodríguez, C., Taylor, G. B., Zavala, R. T., et al. 2006, *Astrophys. J.*, 646, 49

Rossi, E. M., Lodato, G., Armitage, P. J., Pringle, J. E., & King, A. R. 2010, *Mon. Not. R. Astron. Soc.*, 401, 2021

Sądowski, A., Narayan, R., Tchekhovskoy, A., & Zhu, Y. 2013, *Mon. Not. R. Astron. Soc.*, 429, 3533

Schnittman, J. D. 2013, *Classical and Quantum Gravity*, 30, 244007

Schnittman, J. D., Krolik, J. H., & Hawley, J. F. 2006, *Astrophys. J.*, 651, 1031

Schnittman, J. D., Krolik, J. H., & Noble, S. C. 2013, *Astrophys. J.*, 769, 156

Sesana, A., Haardt, F., Madau, P., & Volonteri, M. 2004, *Astrophys. J.*, 611, 623

Stepney, S. & Guilbert, P. W. 1983, *Mon. Not. R. Astron. Soc.*, 204, 1269

Straub, O., Vincent, F. H., Abramowicz, M. A., Gourgoulhon, E., & Paumard, T. 2012, *Astron. Astrophys.*, 543, A83

Teyssier, R. 2002, *Astron. Astrophys.*, 385, 337

Toro, E. F. 1999, *Riemann Solvers and Numerical Methods for Fluid Dynamics* (Springer-Verlag)

Tsokaros, A., Mundim, B. C., Galeazzi, F., Rezzolla, L., & Uryū, K. 2016, *Phys. Rev. D*, 94, 044049

Čada, M. & Torrilhon, M. 2009, *Journal of Computational Physics*, 228, 4118

Valtonen, M. J., Lehto, H. J., Nilsson, K., et al. 2008, *Nature*, 452, 851

van der Holst, B., Keppens, R., & Meliani, Z. 2008, *Computer Physics Communications*, 179, 617

van Meter, J. R., Wise, J. H., Miller, M. C., et al. 2010, *Astrophys. J. Lett.*, 711, L89

Vincent, F. H., Paumard, T., Gourgoulhon, E., & Perrin, G. 2011, *Classical and Quantum Gravity*, 28, 225011

Volonteri, M. 2007, *Astrophys. J. Lett.*, 663, L5

White, C. J. & Stone, J. M. 2015, *ArXiv e-prints [arXiv:1511.00943]*

Younsi, Z. & Wu, K. 2015, *Mon. Not. R. Astron. Soc.*, 454, 3283

Younsi, Z., Wu, K., & Fuerst, S. V. 2012, *Astron. Astrophys.*, 545, A13

Zanotti, O. 2012, *New Astronomy*, 17, 331

Zanotti, O. & Dumbser, M. 2015, *Computer Physics Communications*, 188, 110

Zanotti, O., Fambri, F., & Dumbser, M. 2015, *Mon. Not. R. Astron. Soc.*, 452, 3010

Zanotti, O., Rezzolla, L., Del Zanna, L., & Palenzuela, C. 2010, *Astron. Astrophys.*, 523, A8+

Zanotti, O., Rezzolla, L., & Font, J. A. 2003, *Mon. Not. Roy. Soc.*, 341, 832

Zhang, W. & MacFadyen, A. 2006, *The Astrophysical Journal Supplement Series*, 164, 255

Appendix A: Convergence tests

We have measured the order convergence of the code by studying the norms of the “errors”. Since we employed a finite-volume scheme, the values of the conserved variables u_h at each cell are spatial volume averages of the numerical solution computed at a resolution with cell width h . When an exact solution u exists to the problem at hand, a natural way to quantify the errors is by comparing the computed value of a quantity at each cell with the volume average of the exact solution in the same cell \bar{u} . Formally, the L_1 norm of the error ϵ_h for a resolution h is computed as

$$\|\epsilon_h\|_1 = \|\bar{u}_h - \bar{u}\|_1 \quad (\text{A.1})$$

$$= \sum_{i,j,k} \left| \bar{u}_h - \frac{1}{V_{i,j,k}} \int_{V_{i,j,k}} u \sqrt{\gamma} dx^1 dx^2 dx^3 \right| \quad (\text{A.2})$$

$$= \sum_{i,j,k} |\bar{u}_h - \bar{u}|, \quad (\text{A.3})$$

where $V_{i,j,k}$ is the proper volume of the cell i, j, k , and the sum is performed on all the cells, except those containing the atmosphere, as they are not genuine solutions of the system of partial differential equations. Equivalent relations can be given for the L_∞ norm that keeps track of the maximum error in the domain. The numerical order of convergence of the simulation is then calculated as (see, e.g. [Rezzolla & Zanotti 2013](#))

$$\tilde{p} = \log \left(\frac{\|\epsilon_h\|}{\|\epsilon_k\|} \right) / \log(h/k). \quad (\text{A.4})$$

If the code is convergent, \tilde{p} must be close to the nominal order of accuracy p of the numerical method, which is here 2 for all cases. To simplify the calculations even more, we adopted the refinement factor $h/k = 2$ in this work. In the case of the Michel accretion problem or of a stationary torus, the exact solution is known, so that the relations above at two resolutions can be employed to calculate the convergence order at any time during the evolution.

On the other hand, in the far more common case in which an exact solution is not known, as is the case for the simulations of recoiling black holes, a self-convergence needs to be performed. This requires three different estimates of the errors and the cancellation of the higher-order terms, so that Eq. (A.4) becomes (see, e.g. [Rezzolla & Zanotti 2013](#))

$$\tilde{p} = \log \frac{\|\bar{u}_1 - \bar{u}_2\|}{\|\bar{u}_2 - \bar{u}_3\|} / \log 2, \quad (\text{A.5})$$

where

$$\bar{u}_{2,i,j,k} = \frac{1}{V_{i,j,k}} \int_{i,j,k} u_2 \sqrt{\gamma} dx^1 dx^2 dx^3, \quad (\text{A.6})$$

and

$$\bar{u}_{3,i,j,k} = \frac{1}{V_{i,j,k}} \int_{i,j,k} u_3 \sqrt{\gamma} dx^1 dx^2 dx^3. \quad (\text{A.7})$$

We note that in the expression above, the indices i, j, k refer to the cells of the simulation with the lowest resolution. It is also important to remark that the time-dependent exponent $\tilde{p} = \tilde{p}(t)$ is a genuine measure of the convergence order of our code and provides a far more severe assessment of the convergence properties than the instantaneous measurement shown in Fig. 7. While in




Fig. A.1. Convergence order for the L_1 norm of the 2D recoiling black hole setup with the kick velocity set to zero. We show the uniform grid (solid) and AMR (Löhner scheme) with tolerances of $\varepsilon_t = 0.1$ (dashed) and $\varepsilon_t = 0.01$ (dash-dotted). In the uniform grid case we used three different resolutions: low (128×128), medium (256×256), and high (512×512). In the AMR case, the simulation data used were uniform (128×128 , level 1), and two and three AMR levels.

this work we have presented both approaches to assess the convergence order, measurements of $\tilde{p} = \tilde{p}(t)$ should accompany any work where the convergence properties of a numerical code are presented (see also Radice et al. 2014; Tsokaros et al. 2016).

The infrastructure for refining or coarsening that is present in the code greatly simplifies the task of performing the convergence tests. Since at each refinement level the cell widths are halved with respect to those of the previous level, simulations with higher resolution can be obtained by enforcing a higher level. In practice, the volume averages for Eq. (A.5) are computed through coarsening each snapshot of the data to a lower level. Moreover, when comparing the convergence of the simulations using AMR with that of the uniform cases, a simulation with three AMR levels was taken as equivalent to a uniform simulation with the same resolution of the highest AMR level, and the same averages of Eq. (A.5) were then employed for the convergence test.

Figure A.1 shows the evolution of the convergence order \tilde{p} of the stationary torus, and where we compare simulations with uniform grid and those with two AMR realisations employing a Löhner scheme with tolerances $\varepsilon_t = 0.1$ and $\varepsilon_t = 0.01$. In practice we considered the same setup used when considering a recoiling black hole, but then imposed the kick velocity to be zero. In the uniform-grid case, the resolutions used are ($N_r \times N_\phi$): low (128×128), medium (256×256), and high (512×512). In the AMR case, the same low-resolution case (128×128) was employed, so that the medium and high resolutions are achieved by allowing two and three mesh refinements, respectively.

In this test case, where the torus is stationary and the solution is smooth everywhere except for the torus surface, the convergence order settles to ~ 2.2 in the long-time evolution for the uniform-grid case. This is in good agreement with our expected convergence order, since we have here employed Koren's slope limiter (Koren 1993), which has third-order spatial accuracy in the absence of extrema. For the high-tolerance case with $\varepsilon_t = 0.1$, the AMR run displays a convergence index of only $\tilde{p} \approx 1.8$ in the long-time evolution. However, lowering the tolerance to $\varepsilon_t = 0.01$, we recover the higher-convergence order measured in the uniform-grid case.




Fig. A.2. Convergence (L_1 norm) of the 2D recoiling black hole setup with kick velocity set to $v_{\text{recoil}} = 0.001c$. We show the uniform grid (solid) and AMR (Löhner scheme) with tolerances $\varepsilon_t = 0.1$ (dashed) and $\varepsilon_t = 0.005$ (dash-dotted). In the uniform grid case we used three different resolutions: low (256×128), medium (512×256), and high (1024×512). In the AMR case, the simulation data used were uniform (256×128 , level one), AMR with two levels, and AMR with three levels.

Figure A.2 shows the corresponding convergence order when a kick velocity of $v_R = 10^{-3}c$ is used. Again, we show a uniform case and two AMR cases with tolerances $\varepsilon_t = 0.1$ and $\varepsilon_t = 0.005$. In the uniform-grid case, three different resolutions are employed, with $N_r \times N_\phi = 256 \times 128$ (low), 512×256 (medium), and 1024×512 (high); equivalent AMR realisations are generated allowing one, two, and three mesh refinements.

As clearly shown in Fig. A.2, the convergence order remains higher than 2 in the early stages of the simulation, when the black hole has not yet interacted with the torus matter and the spiral shocks have not yet developed. Most of the simulation region is smooth, hence yielding a high convergence order. In the ensuing stage, the strong shock has developed in the accreting disc and leads to a deterioration of the convergence order, which decreases to being ~ 1 , as expected from Godunov's theorem (see e.g. Rezzolla & Zanotti 2013). After this stage, the spiral shock expands and weakens, and the convergence order increases as the simulation progresses. At the end of the simulation, the convergence order has recovered its stationary value of ~ 2 . The AMR simulations show a similar trend and, in particular, the low-tolerance case is very close to the uniform-grid case at all times.

Finally, the convergence results for the 3D recoiling black hole simulations are shown in Fig. A.3. As seen in the 2D case, the convergence order remains higher than ~ 2 in the early stages of the simulation as most of the solution is smooth. After $t = 10,000 M$, the convergence order gradually decreases, but remains higher than for the 2D case. This is a consequence of the slightly different dynamics of the 3D case. Even though a large spiral shock develops in 3D at around $t = 1,000 M$, the actual accretion onto the black hole, which is responsible for the formation of much of the shock structure seen in 2D, starts only much later at $t \approx 19,500 M$ reported here. As a result, at the end of the simulation, \tilde{p} approaches ~ 1 . The evolution of the convergence order \tilde{p} was also computed for the uniform and AMR simulations, in the same way as described for the 2D case. Clearly, the evolution \tilde{p} for the AMR simulations closely follows that of the uniform cases.




Fig. A.3. Numerical order of accuracy of the 3D recoiling black hole simulation calculated from the L_1 norm using grid (uniform) and AMR (dashed) cases. In the uniform grid case, we used three different resolutions, low ($128 \times 64 \times 16$), medium ($256 \times 128 \times 32$), and high ($512 \times 256 \times 64$). In the AMR runs, tolerances of $\varepsilon_t = 0.1$ and $\varepsilon_t = 0.3$ are set to trigger the second and third refinement levels, respectively. The simulation data used were uniform ($128 \times 64 \times 16$, level 1), AMR with two levels, and AMR with three levels.